Computerising Diagnosis:
Minds, Medicine, and Machines in Twentieth-Century America



Andrew Lea
History of Science and Medicine
Balliol College
University of Oxford

A thesis presented for the degree of
Doctor of Philosophy
2019

Name: Andrew Lea

College: Balliol College

Dissertation: Computerising Diagnosis:
Minds, Medicine, and Machines in Twentieth-Century America

Degree: Doctor of Philosophy (History of Science and Medicine)

Term: Michaelmas Term 2019

Abstract

In the mid-twentieth century, physicians, engineers, mathematicians, philosophers, and others started to think about the ways in which a new technology might be applied to an old problem. These interdisciplinary teams designed computer programmes that aimed to automate one of the most fundamental tasks of medicine—diagnosis. This dissertation offers a history of computerising diagnosis through the prism of three such efforts: the Medical Data Screen, a computerised method of capturing the health of 'the total patient'; HEME, a computer programme created to aid in the diagnosis of haematologic diseases; and MYCIN, an early expert system developed to assist in the diagnosis and treatment of microbial infections. Drawing on these case studies, this project examines early efforts to computerise medical diagnosis and decision making. It explores how these efforts built upon, interacted with, and produced certain professional tensions, disease constructions, personal identities, cultural ideals, economic interests, and material practices.

This dissertation supports a number of conclusions. First, it demonstrates that efforts to computerise medical diagnosis and decision making raised philosophical and moral questions that lie at the heart of medical practice. As physicians and engineers worked to computerise diagnosis, they were forced to reckon with fundamental epistemological and ontological questions—questions about the nature of clinical reasoning and the definitions of disease. In creating computerised diagnostic systems, early developers articulated new ideas about how physicians think and how disease categories are defined. Second, this dissertation argues that the introduction of computers into medical diagnosis was both the cause *and* the effect of new ways of collecting, storing, and using medical information. Third, this dissertation highlights the complexities inherent to medicine and medical diagnosis. Repeatedly early developers of computerised diagnostic systems ran up against unanticipated epistemological, social, and moral challenges.

Extended Abstract

In the mid-twentieth century, physicians, engineers, mathematicians, philosophers, and others started to think about the ways in which a new technology might be applied to an old problem. These interdisciplinary teams designed computer programmes that aimed to automate one of the most fundamental tasks of medicine—diagnosis. This dissertation offers a history of computerising diagnosis through the prism of three such efforts: the Medical Data Screen, a computerised method of capturing the health of 'the total patient'; HEME, a computer programme created to aid in the diagnosis of haematologic diseases; and MYCIN, an early expert system developed at Stanford to assist in the diagnosis and treatment of microbial infections. Drawing on these case studies, this project examines early efforts to computerise medical diagnosis and decision making. It explores how these efforts built upon, interacted with, and produced certain professional tensions, disease constructions, personal identities, cultural ideals, economic interests, and material practices.

This dissertation supports a number of conclusions. First, it argues that efforts to computerise medical diagnosis and decision making raised philosophical and moral questions that lie at the heart of medical practice. As physicians and engineers worked to computerise diagnosis, they were forced to reckon with fundamental epistemological and ontological questions—questions about the nature of clinical reasoning and the definitions of disease. In creating computerised diagnostic systems, early developers articulated new ideas about how physicians think and how disease categories are defined. Sometimes these ideas and categories were spawned by their engagements with computers; other times, pre-existing ideas came to be codified within subsequent computer programmes. Moral questions travelled together with these philosophical questions. The computerisation of diagnosis broached profound moral and professional uncertainties of ceding important medical decisions to a machine.

Second, this dissertation argues that the introduction of computers into medical diagnosis was both the cause *and* the effect of new ways of collecting, displaying, storing, and using medical information. Through the twentieth century, physicians and medical scientists developed new, standardised modes of gathering, manipulating, storing, and reporting medical information. The computer's foray into medical diagnosis grew out of these earlier, paper-based tools and practices. Yet computerisation also created the need for the transformation of medical information into forms that were suitable to the material conditions of the machine.

Third, this dissertation highlights the complexities inherent to medicine and medical diagnosis. Over the second half of the twentieth century, physicians and engineers pursued strategies of all kinds to domesticate the messy work of diagnosis—to make it amenable to computerisation. In their attempts to computerise medical diagnosis, these actors repeatedly ran up against, and underestimated, medicine's complexities. They made assumptions that did not hold. They raised questions that had no quick or easy answers. They overlooked ethical and moral dimensions of medical decision making. The efforts tracked in this dissertation raised fascinating questions, stirred heated discussions, and animated a wide spread of hopes and fears. But ultimately, they all failed. To be sure, the reasons for this failure were many, pertaining to context-dependent technological, organisational, social, and conceptual limitations and shortcomings. Yet, alongside these varied external sources of failure was a source internal to medicine itself: the complexity of diagnosis, the intricacy of patient care.

The dissertation proceeds in three sections, each comprising two chapters and each organised around a particular case in the history of computerising diagnosis. Part 1, 'Paper Infrastructures', traces the creation of the Medical Data Screen (MDS), a computerised method of generating differential diagnoses. Though a novel digital method, the MDS in fact had its origins in a paper-based analogue technology: the medical history questionnaire. In 1947, the Cornell psychiatrist Keeve Brodman and a handful of colleagues began developing what would become one of the most widely used health questionnaires of its time—the Cornell Medical Index (CMI). A rigidly standardised form, the CMI presented 195 yes-no questions designed to capture the health status of 'the total patient'. Chapter 1 traces the development of the CMI: its motivating aims, its surprising uses, and its material grounding.

Over the following decades, Brodman's project of standardising medical history taking gradually evolved into a project of mathematising and computerising diagnosis: out of the CMI grew the Medical Data Screen (MDS), one of the first computerised methods of deriving diagnoses from patient data. Chapter 2 explores the process of converting an analogue paper technology into a digital computerised one. The computerised MDS evolved from, and rested upon, paper data practices associated with questionnaires and surveys. Both tools were motivated by similar hopes. In an era of overstretched and underresourced physicians, Brodman and his colleagues believed that these technologies would not only extend the reach of medical personnel to rural and underserved areas but also relinquish physicians from the monotonous aspects of medical practice, giving them more time to focus on the 'human' aspects of care. The CMI and MDS also stirred similar debates and tensions. The introduction of decision support technologies into medical history taking and diagnosis fuelled concerns around professional deskilling, interpersonal dehumanisation, and substandard care. Part 1 of this dissertation thus foregrounds the important continuities between computer-based digital practices in medicine and paper-based analogue ones.

Part 2, 'Recoding Disease', examines the role of the computer—and material culture more generally—as an agent of disease definition. The section explores this theme through the lens of a thirty-five-year effort to computerise the diagnosis of haematologic diseases. In 1957, a group of researchers based at the Rockefeller Institute gathered at the Radio Corporation of America (RCA) facility in Camden, New Jersey to witness the RCA's Bizmac computer render, from inputs of patient data, diagnoses of haematologic diseases. Chapter 3 follows the development and evolution of this haematology programme. It shows how physicians and engineers, as they encountered novel digital media, began to represent medical information and define diseases in new ways. In exploring the computer's role as an agent of disease definition, this chapter suggests that the process of disease (re)definition is not just a social or intellectual process. It is a material process as well, shaped profoundly by the material media with which physicians work. Definitions of disease have a body—they are incarnate in stuff.

As interdisciplinary teams worked to automate medical diagnosis, they repeatedly bumped up against fundamental epistemological questions. How do doctors think? What is the nature of human reasoning? Are human modes of reasoning amenable to computerisation? If they are, *should* computers aim to mimic them? In what ways do physicians come to define and classify disease categories? Are these categories even necessary? While following the haematology programme into its final decades, Chapter 4 explores how computers both generated and incorporated ideas about the medical mind and its inner workings.

Part 3, ‘Black Box Medicine’, uses the case of MYCIN to examine key moral tensions attendant to the computer’s introduction into medical diagnosis. In the early 1970s, the young MD-PhD student Edward H. Shortliffe started work on his dissertation within Stanford University’s Department of Computer Science. His PhD thesis, completed in 1975, involved the development of a computer programme called MYCIN, an expert system that provided physicians with advice concerning the diagnosis and treatment of bacterial infections. MYCIN immediately attracted wide attention: it matured into a much larger research enterprise at Stanford and spawned larger conversations about whether and how to apply computers to the problem of medical diagnosis. Chapter 5 recounts this early effort to computerise medical diagnosis and decision making. It demonstrates that questions about trust and transparency animated the development of, and responses to, the MYCIN project. How could a physician using the system know that its advice was accurate and trustworthy? What did the developers of MYCIN do to make the system’s reasoning comprehensible to its human users? The MYCIN team largely pursued technological solutions to the problems of trust and transparency, creating novel technologies, programmes, and interfaces that purportedly allowed MYCIN to ‘explain itself’.

As Chapter 6 argues, these issues around trust and transparency were entangled with concerns over authority, identity, and responsibility. To cede important treatment and diagnostic decisions to machine—to allow computers to carry out a fundamental task of medicine—raised questions about what it meant to be a physician as well as questions about the physician’s responsibility and his ownership over medical decisions. Despite all that the MYCIN did to render their system’s reasoning transparent, physicians worried about lingering opacities of a computerised system—about the biases that, as one actor put, may still be ‘hidden in the code’. The technological fixes to issues of transparency thus fell short: they could not cleanly resolve the larger professional, regulatory, moral, and epistemological challenges raised by computerised medicine.

Acknowledgements

This dissertation reflects the intellectual generosity of numerous individuals spread across several institutions. I have been fortunate to have Oxford's Wellcome Unit for the History of Medicine as my intellectual home base. My supervisor Mark Harrison shaped this dissertation every step of the way, providing invaluable advice about my driving questions, sources, and arguments. He has been an inexhaustible source of insight and encouragement. At Oxford I also benefitted from conversations with Erica Charters, Pietro Corsi, Robert Iliffe, Michael Joseph, Claas Kirchhelle, Laura Lamont, Sloan Mahone, Margaret Pelling, Steve Server, Katherine Warren, and Steve Woolgar. Belinda Clark's friendship and administrative support kept me both nurtured and productive.

Many others outside of Oxford contributed to this project. At the Max Planck Institute for the History of Science I benefited from a rich and vibrant institutional culture as a Predoctoral Fellow in Department II. Robert Aronowitz, Lorraine Daston, Cathy Gere, J. Andrew Mendelsohn, and David Sepkoski all provided incisive feedback and advice as I was beginning the process of writing.

Even as a doctoral student, I continued to learn from mentors I had been lucky enough to meet as an undergraduate at Harvard University. David S. Jones and Scott Podolsky read and commented on large portions of this dissertation. Their enthusiasm for my research sustained me, and their advice improved this work immeasurably.

I am particularly indebted to Jeremy A. Greene, who read and commented on multiple versions of this dissertation. Our conversations—about history, computing, and clinical medicine—have helped form many of the ideas presented in the dissertation as well as my broader sense of what it means to pursue a career that bridges historical scholarship and clinical practice. This dissertation would not be what it is were it not for his extraordinary intellectual energy and generosity.

Finally, this dissertation would not have been possible without generous funding from The Rhodes Trust, the Max Planck Institute for the History of Science, and the Hagley Museum and Library. Archivists and librarians at many institutions went above and beyond the call of duty in helping me locate critical materials for this project.

Above all, I am grateful to my family for their love and support.

Contents

Introduction: ‘Now Dr. Robot Enters the Scene’	1
Part I: Paper Infrastructures	**24**
Chapter I: Indexing the World	25
Chapter II: The Statistical Patient	60
Part II: Recoding Disease	**128**
Chapter III: ‘The Standardized Language of the Machine’	129
Chapter IV: The Medical Mind	162
Part III: Black Box Medicine	**202**
Chapter V: MYCIN Explains Itself	203
Chapter VI: ‘Hidden in the Code’	254
Conclusion and Epilogue: Looking Forward, Looking Back	305
Figures	327
Bibliography	348

Abbreviations

ACME	Advanced Computer for Medical Research
AI	Artificial Intelligence
AIL	Airborne Instruments Laboratory
AIM	Artificial Intelligence in Medicine
AJM	American Journal of Medicine
AMA	American Medical Association
AMPAC	American Medical Political Action Committee
APA	Automated Physician's Assistant
APHAS	Automated Patient History Acquisition System
BBN	Bolt, Beranek and Newman
BMJ	British Medical Journal
CASNET	Causal Associational Network
CDSS	Clinical Decision Support System
CF	Certainty Factor
CMI	Cornell Medical Index
CNS	Central Nervous Systems
CPC	Clinicopathological Conference/Computer Policy Committee
CSI	Cornell Service/Selectee Index
DARPA	Defense Advanced Research Projects Agency
DENDRAL	Dendritic Algorithm
ECG	Electrocardiogram
EHR	Electronic Health Record
EMR	Electronic Medical Record
ENIAC	Electronic Numerical Integrator and Computer
FDA	U.S. Food and Drug Administration
GDPR	General Data Protection Regulation
GIGO	Garbage In, Garbage Out
HELP	Health Evaluation through Logical Processing
HEW	U.S. Department of Health, Education and Welfare
HPP	Heuristic Programming Project
IBM	International Business Machines
IRL	Instrumentation Research Laboratory
JAMA	Journal of the American Medical Association
JPL	Jet Propulsion Laboratory
LCS	Laboratory of Computer Science
LDS	Latter-day Saints
LHASA	Logic and Heuristics Applied to Synthetic Analysis
Lisp	List-Processing Language
LT	Logic Theorist
MDS	Medical Data Screen
MGH	Massachusetts General Hospital
MIT	Massachusetts Institute of Technology
MMPI	Minnesota Multiphasic Personality Inventory
MSTP	Medical Science Training Program
NASA	National Aeronautics and Space Administration
NBS	National Bureau of Standards

NEJM	New England Journal of Medicine
NIH	National Institutes of Health
NYU	New York University
OR	Operations Research
PACS	Picture Archiving and Communication System
PHS	Public Health Service
QA	Question-Answer
RCA	Radio Corporation of America
RCT	Randomised Controlled Trial
RMP	Regional Medical Program
ROS	Review of Systems
SCC	Stanford Computation Center
SGO	Surgeon General's Office
SEAC	Standards Eastern Automatic Computer
SMDM	Society for Medical Decision Making
STEM	Science, Technology, Engineering, and Mathematics
STS	Science and Technology Studies
SUMEX-AIM	Stanford University Medical Experimental Computer for Artificial Intelligence in Medicine
UCSF	University of California, San Francisco
UGDP	University Group Diabetes Program
XAI	Explainable AI

INTRODUCTION

'Now Dr. Robot Enters the Scene'

On 8 December 1956, millions of Americans tuned in to watch an episode of the popular television series *Science Fiction Theatre.* The episode, 'Dr. Robot', followed a team of researchers working on a secret scientific project known as 'Operation Polyglot'. The goal of the project, as its named suggested, was to develop a computer that would serve as a mechanical translator of languages.[1] After a long day of work, Edgar Barnes, the project lead, found evidence that someone had been tampering with the machine: wires, leading to certain memory units of the uncompleted machine, had been soldered to the computer terminal. Cold War tensions high, Barnes's mind leapt immediately to the threat of espionage.

With the help of a security officer, Barnes began to quietly investigate the project's chief engineer, mathematician, and linguist—the only other personnel with unfettered access to the machine. In the course of monitoring these men, Barnes discovered that the chief mathematician, Fred Lopert, had indeed been tampering with the machine—but not for the suspected nefarious reason. He was merely using the machine to decide whether his wife, afflicted with a life-threatening heart condition, should continue with the present course of pharmaceutical treatment or undergo invasive surgery. 'I was hoping', Lopert explained to Barnes, 'the machine could

[1] Machine translation (MT) was an active, albeit fledgling, area of research during this period. Much of this research was animated by fears that too few American scientists and government officials were able to monitor or comprehend the scientific developments coming out of the Soviet Union. Michael Gordin, *Scientific Babel: How Science Was Done Before and After Global English* (London and Chicago: University of Chicago Press, 2015), pp. 213-66.

analyze all the facts of her case and decide which method had the best chance of saving her'.

Barnes, both sympathetic to Lopert's situation and stimulated by the intellectual challenge, allowed his associate to carry out the project. Lopert's undertaking proved so successful that the foundation sponsoring Operation Polyglot asked the researchers to permanently convert the computer 'into a master medical consultant'. Asked how the new machine would work, Barnes replied: 'Well, a physician with a difficult diagnostic problem would fill out a form describing the problem and giving all the facts in the case. Then, the machine would compare it with the information gathered from similar cases all over the world. And then it would tell you what your chances are on the various methods of treatment'.

During the mid-twentieth century, ideas about applying computers to the problem of medical diagnosis were not confined to television—nor to science fiction. As countless Americans enjoyed 'Dr. Robot' on their television sets, a small but growing group of physicians, medical researchers, psychologists, mathematicians, and computer scientists were beginning to explore the question of how to computerise medical diagnosis and decision making.[2] Indeed, the episode concluded by suggesting that the developments depicted in the show were just around the corner—if not already under way: 'When this story was conceived, we believed that it was fiction. But science marches on, and at the last cybernetics congress a machine was displayed that diagnoses disease. And once again the dream of a science fiction writer becomes a

[2] The 26 May issue of *The Billboard* puts the March 1956 rating of *Science Fiction Theatre* at 12.6. With an estimated thirty-five million television-owning U.S. households in 1956, one can extrapolate from these data that the programme's viewership was likely in the millions. 'The Billboard Scoreboard: Syndicated Film Programs', *The Billboard,* 26 May 1956, p. 9. For a detailed historical discussion of the programme, see Martin Grams, Jr., *Science Fiction Theatre: A History of the Television Program, 1955-1957* (Albany, GA: BearManor Media, 2011).

reality'. By 1963, the *New York Times* had also declared the passage of the computer diagnostician from the television screen to the hospital ward, running the headline, 'Now Dr. Robot Enters the Scene'.[3]

This dissertation offers a history of computerising medical diagnosis through the prism of three case studies. In 1947, the Cornell psychiatrist Keeve Brodman and a handful of colleagues began developing what would become one of the most widely used health questionnaires of its time—the Cornell Medical Index (CMI). A rigidly standardised form, the CMI presented 195 yes-no questions designed to capture the health status of 'the total patient'. Over the following decades, Brodman's project of standardising medical history taking gradually evolved into a project of mathematising and computerising diagnosis: out of the CMI grew the Medical Data Screen (MDS), an early computerised method of deriving diagnoses from patient data. At the same time Brodman was beginning to work towards the MDS, another research team, headed by the television pioneer Vladimir Zworykin, was developing a computer programme that they hoped would make accurate diagnoses in the field of haematology. By the early 1970s, other researchers had begun to pursue self-avowedly 'artificial intelligence' (AI) approaches to computerising diagnosis. At Stanford, the young MD-PhD student Edward H. Shortliffe developed a computer programme called MYCIN, an early expert system that aimed to provide advice concerning the diagnosis and treatment of bacterial infections.

While far from exhaustive, these efforts were among the earliest and most widely discussed attempts to bring medical diagnosis under the remit of the digital computer. They all generated considerable discussion about whether, how, and why to apply computers to arguably the most foundational, intimate, and cerebral task of

[3] Louis Lasagna, 'Now Dr. Robot Enters the Scene', *New York Time*, 16 June 1963.

medicine—diagnosis. Drawing on these case studies, this dissertation explores how attempts to computerise diagnosis built upon, interacted with, and produced certain professional tensions, economic interests, disease constructions, cultural ideals, and material practices. Debates about how and whether to computerise diagnosis were often animated by larger debates about the nature of diagnosis and medical reasoning, the definitions of disease, and the authority as well as the identity of physicians and their patients.

This dissertation supports a number of arguments and conclusions. First, it argues that efforts to computerise medical diagnosis and decision making raised philosophical and moral questions that lie at the heart of medical practice. As physicians and engineers worked to digitise diagnosis, they were forced to reckon with fundamental epistemological and ontological questions—questions about the nature of clinical reasoning and the definitions of disease. In creating computerised diagnostic systems, early developers articulated new ideas about how physicians think and how disease categories are defined. Sometimes these ideas and categories were spawned by their engagements with computers; other times, pre-existing ideas came to be codified within subsequent computer programmes. Moral questions travelled together with these philosophical questions. The computerisation of diagnosis broached profound moral and professional uncertainties of ceding important medical decisions to a machine.

Second, this dissertation argues that the introduction of computers into medical diagnosis was both the cause *and* the effect of new ways of collecting, displaying, storing, and using medical information. Through the twentieth century, physicians and medical scientists developed new, standardised modes of gathering, manipulating, storing, and reporting medical information. The computer's foray into medical diagnosis grew out of these earlier, paper-based tools and practices. Yet

computerisation also created the need for the transformation of medical information into forms that were suitable to the material conditions of the machine. This was the process, to use a phrase of one of this dissertation's actors, of translating medical information into 'the standardized language of the machine'. The computer's emergence within medicine, moreover, created new cognitive models about what it meant for medical information to be 'truly standardised'.

Finally, this dissertation highlights the complexities inherent to medicine and medical diagnosis. Over the second half of the twentieth century, physicians and engineers pursued strategies of all kinds to domesticate the messy work of diagnosis—to make it amenable to computerisation. In their attempts to computerise medical diagnosis, these actors repeatedly ran up against, and underestimated, medicine's complexities. They made assumptions that did not hold. They raised questions that had no quick or easy answers. They overlooked ethical and moral dimensions of medical decision making. The efforts tracked in this dissertation raised fascinating questions, stirred heated discussions, and animated a wide spread of hopes and fears. But ultimately, they all failed. To be sure, the reasons for this failure were many, pertaining to context-dependent technological, organisational, social, and conceptual limitations and shortcomings. Yet, alongside these varied external sources of failure was a source internal to medicine itself: the complexity of diagnosis, the intricacy of patient care.

HISTORIES OF COMPUTERISED MEDICINE

This dissertation sits at the interface of multiple literatures. Computers were established as subjects of sustained historical inquiry in the final years of the twentieth century, with the formative work of many historians including Paul Edwards, Martin

Campbell-Kelly, and William Aspray.[4] Soon thereafter, historical perspectives on artificial intelligence and machine learning also began to find their way into the scholarly literature.[5] Despite this proliferation of scholarship on the history of computing, the introduction of computing into medicine and medical diagnosis remains an underexplored topic, eclipsed by an interest in the application of computers to other domains, including mathematics and military strategy.[6] Just as medicine has been an underexamined topic among historians of computing, computing has long been an underexamined topic among historians of medicine. Medical technology has long been a rich and productive focus within the history of medicine.[7] But until recently historians

---

[4] Martin Campbell-Kelly and William Aspray, *Computer: A History of the Information Machine* (New York: Basic Books, 1996); Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America* (Cambridge, MA: MIT Press, 1997); Agatha Hughes and Thomas Parke Hughes, *Systems, Experts, and Computers: The Systems Approach in Management and Engineering, World War II and After* (Cambridge, MA: MIT Press, 2000). Among the more recent additions to the literature are Atsushi Akera, *Calculating a Natural World: Scientists, Engineers, and Computers During the Rise of U.S. Cold War Research* (Cambridge, MA: MIT Press, 2008); Michael Mahoney, 'The Histories of Computing(s)', *Interdisciplinary Science Reviews* 30, no. 2 (2005): 119-135.

[5] See, e.g., Raymond Kurzweil, *The Age of Intelligent Machines* (Cambridge, MA: MIT Press, 1990); Nils Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Cambridge and New York: Cambridge University Press, 2010). Important recent work includes Stephanie Dick, 'AfterMath: The Work of Proof in the Age of Human-Machine Collaboration', *Isis* 102, no. 3 (2011): 494-505; Stephanie Dick, 'Of Models and Machines: Implementing Bounded Rationality', *Isis* 106, no. 3 (2015): 623-634; Matthew L. Jones, 'How We Became Instrumentalists (Again): Data Positivism since World War II', *Historical Studies in the Natural Sciences* 48, no. 5 (2018): 673-684. On the challenges of writing the history of machine learning, see Aaron Plasek, 'On the Cruelty of Really Writing a History of Machine Learning', *IEEE Annals of the History of Computing* 38, no. 4 (2016): 6-8.

[6] There is also a flourishing literature on the history of (big) data. See, e.g., Elena Aronova, Christine von Oertzen, and David Sepkoski, 'Introduction: Historicizing Big Data', *Osiris* 32, no. 1 (2017): 1-17; Lorraine Daston (ed.), *Science in the Archive: Pasts, Presents, Futures* (Chicago and London: University of Chicago Press, 2017); Rebecca Lemov, *Database of Dreams: The Lost Quest to Catalog Humanity* (New Haven: Yale University Press, 2015).

[7] For helpful, albeit somewhat outdated, reviews of the literature, see Jennifer Stanton, 'Making Sense of Technologies in Medicine', *Social History of Medicine* 12 (1999): 437-448; Harry M. Marks, 'Medical Technologies: Social Contexts and Consequences', in W. F. Bynum and Roy Porter (eds.), *Companion Encyclopedia of*

of the life sciences, medicine, and medical technology have paid insufficient attention to the emergence of digital electronic computing in medical practice.[8] This dissertation bridges this divide.

A string of recent works has also started to disrupt these disciplinary margins, with historians of computing identifying the life sciences as a site of productive historical inquiry and historians of the life sciences attending increasingly to the role of computation in their chosen area of interest. In, for example, his book *Life Out of Sequence*, Hallam Stevens examines the scope, availability, and shape of biological data as a way of charting the transformations to biomedical research wrought by the integration of information technology.[9] In a more historically-oriented work, the historian Joseph November sketches out the origins of biomedical computing in the

---

the History of Medicine, Vol. 2, (New York: Routledge, 1993), 1592-1618. For some specific examples, see, e.g., Thomas Schlich and Christopher Crenner (eds.), *Technological Change in Modern Surgery: Historical Perspectives on Innovation* (Melton: University of Rochester Press, 2017); Carsten Timmermann and Julie Anderson (eds.), *Devices and Designs: Medical Technologies in Historical Perspective* (New York: Palgrave Macmillan, 2006); Keith Wailoo, *Drawing Blood: Technology and Disease Identity in Twentieth-Century America* (Baltimore: Johns Hopkins University Press, 1997).

[8] Of course there are notable exceptions. Warwick Anderson, 'The Reasoning of the Strongest: The Polemics of Skill and Science in Medical Diagnosis', *Social Studies of Science* 22, no. 4 (1992): 653-684; Brian Dolan and Allison Tillack, 'Pixels, Patterns and Problems of Vision: The Adaptation of Computer-Aided Diagnosis for Mammography in Radiological Practice in the US', *History of Science* 48 (2010): 227–249; Stanley Joel Reiser, *Medicine and the Reign of Technology* (Cambridge and New York: Cambridge University Press, 1978); Stanley Joel Reiser, *Technological Medicine: The Changing World of Doctors and Patients* (Cambridge and New York: Cambridge University Press, 2009).

[9] Hallam Stevens, *Life Out of Sequence: A Data-Driven History of Bioinformatics* (Chicago: University of Chicago Press, 2013). Other works that explore the application of computation to biomedical research include Lily Kay, *Who Wrote the Book of Life? A History of the Genetic Code* (Stanford: Stanford University Press, 2000); C. Stewart Gillmor, *Fred Terman at Stanford: Building a Discipline, a University, and Silicon Valley* (Stanford: Stanford University Press, 2004); Bruno J. Strasser, *Collecting Experiments: Making Big Data Biology* (Chicago: University of Chicago Press, 2019).

United States through the 1950s and 1960s.[10] Beyond his many insights concerning the relationship of biomedical computing to research agendas, governmental funding, and health infrastructures, November demonstrates that biomedical research and computer technology were mutually constitutive: just as the objects of medicine and biology evolved to become 'systems that computers could indeed describe', computers also changed as a result of their interactions with biomedical research.[11]

Historians, though, have not been the only—or even the primary—narrators of the history of electronic computation in a medical context; since the late twentieth century, key historic agents in the founding of medical informatics as an established field have served as authors of their own contributions and experiences. The edited volume *The History of Medical Informatics*, now in its second edition, synthesises the contributions of pioneers in the development of medical informatics to assemble a story of the field's evolution that is meticulous in detail and rich in insight.[12] This historical account is supplemented by a number of other works that range from historical monograph to personal memoire.[13] Collectively, these texts cover the many domains of medicine that have been variously shaped by computers, including imaging systems, electronic health records, medical information systems, and, to a limited extent, clinical decision support systems. As full of a picture as these texts provide, internally-

---

[10] Joseph November, *Biomedical Computing: Digitizing Life in the United States* (Baltimore: Johns Hopkins University Press, 2012).
[11] November, *Biomedical Computing*, p. 7.
[12] Morris Collen (ed.), *A History of Medical Informatics in the United States, 1950 to 1990* (Indianapolis, IN: American Medical Informatics Association, 1995); Morris Collen and Marion Ball (eds.) *The History of Medical Informatics in the United States, Second Edition* (London: Springer, 2015).
[13] Bruce Blum and Karen Duncan (eds.), *A History of Medical Informatics: Proceedings of Association for Computing Machinery Conference on History of Medical Informatics* (New York: ACM Press, 1990); Severo Ornstein, *Computing in the Middle Ages: A View from the Trenches, 1955-1983* (Bloomington: 1stBooks, 2002).

produced accounts often espouse triumphalist and teleological narratives of their discipline.

This dissertation departs from these existing works in a number of ways. Thematically, where the large bulk of historical work at the interface of biomedicine and computing has gravitated towards the computer's role in biomedical research, this dissertation engages with computation in clinical practice and diagnostics. This study, in other words, acts as a counterweight to the historiographical privileging of the computer's place at the lab bench over its place at the bedside.[14] Conceptually, this dissertation's critical engagement with questions concerning disease ontology, medical epistemology, and professional identity goes beyond a reconstruction of the order in which key events unfolded. This dissertation's interest, then, is less in establishing a full picture of *what happened* with respect to early computerised systems in medicine than in untangling and identifying the various preconditions and ramifications of what happened.

In posing questions of this conceptual orientation, this dissertation is beholden to the prodigious literatures coming out of science and technology studies (STS). In particular, the literature on how new health technologies affect both medical practices and human lives offers particularly valuable analytic frameworks for understanding the historical significance of computerised diagnosis.[15] In fact, a number of STS scholars,

---

[14] Although, as November points out, the computer's use in biological research and its use in clinical medicine were often lumped together under the common rubric of 'biomedical computing', actors during the mid-twentieth century recognised that the 'bio' and the 'medical' sides to computing each came with a unique set of questions, frameworks, and challenges. November, *Biomedical Computing*, p. 173.

[15] Nik Brown and Andrew Webster, *New Medical Technologies and Society: Reordering Life* (Cambridge and Maldan: Polity, 2004); Harry Collins and Trevor Pinch, *Dr. Golem: How to Think about Medicine* (Chicago: University of Chicago Press, 2005); Alex Faulkner, *Medical Technology into Healthcare and Society: A Sociology of Devices, Innovation and Governance* (New York and London: Palgrave Macmillan, 2009); Christian Heath, Paul Luff, and Marcus Sanchez Svensson,

Marc Berg chief among them, have analysed the implementation of computerised medicine using the field's tools and frameworks.[16]

In a notable study on emerging telecare technologies, Nelly Oudshoorn reveals how these technologies participate in redefining technical devices, professional identities, and the norms and geographies of care.[17] Yet, as other threads within the STS scholarship make clear, such technology-driven transformations tend to be characterised by multiplicity and contingency: medical technologies have no stable or essential nature—they instead manifest differently as they are enacted in different locales and at different network nodes.[18] In this vein, Annemarie Mol has shown in her

---

'Technology and Medical Practice', *Sociology of Health and Illness* 25, no. 3 (2003): 75-96; Ericka Johnson and Boel Berner, *Technology and Medical Practice: Bloods, Guts and Machines* (Farnham and Burlington: Ashgate Publishing Company, 2010); Pascale Lehoux, *The Problem of Health Technology: Policy Implications for Modern Health Care Systems* (New York and London: Routledge, 2006); Annemarie Mol, *The Logic of Care: Health and the Problem of Patient Choice* (London and New York: Routledge, 2008); Sonja Olin-Lauritzen and Lars-Christer Hydén, *Medical Technologies and the Life World: The Social Construction of Normality* (London and New York: Routledge, 2007); Nelly Oudshoorn and Trevor Pinch, *How Users Matter: The Co-Construction of Users and Technologies* (Cambridge, MA: MIT Press, 2003); Stefan Timmermans and Marc Berg, 'The Practice of Medical Technology', *Sociology of Health and Illness* 25 (2003): 97-114.

[16] Marc Berg's work on decision support technologies in twentieth-century medicine is exceptionally instructive in making sense of the legacy and impact of expert systems. Through the lens of three different examples of decision support, Berg shows how decision support tools restructure medical work. Central to Berg's account is the idea that decision support tools on one hand, and clinical practices and personnel on the other, mutually constitute one another. Just as 'rational aids' are transformed by their users and surroundings, so these aids transform these people and places in a reciprocal fashion, disciplining them 'to a formalism'. Marc Berg, *Rationalizing Medical Work: Decision Support Techniques and Medical Practices* (Cambridge, MA: MIT Press, 1997). See also Diana Forsythe, 'Engineering Knowledge: The Construction of Knowledge in Artificial Intelligence', *Social Studies of Science* 23, no. 3 (1993): 445-477; Diana Forsythe*, Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence* (Stanford: Stanford University Press, 2002); Sarah Nettleton, 'The Emergence of E-Scaped Medicine?', *Sociology* 38, no. 4 (2004): 662-679.

[17] Nelly Oudshoorn, *Telecare Technologies and the Transformation of Healthcare* (Basingstoke: Palgrave Macmillan, 2011).

[18] Marc Berg and Annemarie Mol, *Differences in Medicine: Unraveling Practices, Techniques, and Bodies* (Durham and London: Duke University Press, 1998). Studies that drawn upon actor-network theory (ANT) frequently emphasise the dynamic and

ethnography of atherosclerosis that this seemingly uncomplicated condition is in fact not a single condition at all but an illness of multiple realities—a disease that looks markedly different depending on who one is (e.g., surgeon versus radiologist) and where one is situated (e.g., laboratory versus examination room).[19] Certain medical technologies, including radiological imaging and other diagnostic tools, play a central role in the fracturing of this disease concept. Running through this collection of disparate studies is a pervasive theme: medical technology 'come[s] to feature in *determinate, but unpredictable* ways' in the structuring, practices, and logics of healthcare.[20] Computerised diagnostic systems are no exception. Although it is true that these technologies left distinctive imprints on medical thought and practice, the precise contours of these imprints varied according to an array of situational factors, local contingencies, and historical continuities. By taking a case-studies approach to the history of computerised diagnosis, this dissertation aims to straddle between particularity and breadth—to bring light to general transformations without losing sight of the embeddedness and continuities of medical and technological change.

This dissertation makes a number of contributions to this literature. First, it reconstructs three efforts to computerise diagnosis and decision making that have received scant historical attention. These were not only some of the earliest efforts to

---

effervescent character of medical technologies. Madeline Akrich, 'The De-Scription of Technical Objects', in Wiebe E. Bijker and John Law (eds.), *Shaping Technology / Building Society: Studies in Sociotechnical Change* (Cambridge, MA: MIT Press 1992), 205-224. For an introduction to ANT, see Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network Theory* (Oxford: Oxford University Press, 2005). For some of the work that laid the groundwork for ANT, see Bruno Latour and Stephen Woolgar, *Laboratory Life: The Construction of Scientific Facts* (Beverely Hills: Sage, 1979).

[19] Annemarie Mol, *The Body Multiple: Ontology in Medical Practice* (Durham: Duke University Press, 2002).

[20] Heath, Luff, and Sanchez Svensson, 'Technology and Medical Practice', p. 81. Emphasis added.

computerise diagnosis but also some of the most influential: they were published in leading medical journals, covered by the national press, and discussed at national and international meetings. In reconstructing these efforts, this dissertation brings new archival materials into the historical record while at the same time working to situate and suspend these materials within their historical and intellectual contexts. More generally, this dissertation represents the first monograph-length history of computerising diagnosis—a topic that spurred a remarkable level of intellectual activity, and attracted a great deal of attention and discussion, through the latter half of the twentieth century.

A second contribution of this dissertation is interpretive. The conclusions supported by this dissertation (outlined above) offer new insights into how the computer came to be applied to the problem of medical diagnosis—and the broader meanings and consequences of these applications.

Finally, this dissertation hopes to make a methodological contribution to the literature. Many of the sources and methods of this dissertation will be familiar to the historian of medicine. This dissertation employs a wide range of material to reconstruct early efforts to computerise medical diagnosis. In addition to the published medical, mathematical, and computer science literature, this dissertation uses numerous archival paper collections, consisting of letters, notes, essays, speeches, and much else, to understand and unpack these efforts—their aims, aspirations, and assumptions. Other administrative records—meeting minutes, grant applications, and conference proceedings—supplement these sources, offering insights into how these computerised systems were imagined and developed.

Because this dissertation tracks not just the emergence of early computerised systems but also the responses to them, it also explores various local, national, and

international newspapers and magazines to determine how various groups—physicians, engineers, and lay people—responded to attempts to bring the computer into the arena of medicine. Professional magazines and periodicals, such as *Medical World News* and *The Hospital Physician*, offer useful perspectives from within the medical profession. As is the case with many recent histories of medicine, the patient's perspective is largely absent from this dissertation—largely a consequence of privacy regulations and a dearth of sources at this time.

Yet this dissertation also brings new kinds of sources to bear on topics and questions in the history of medicine. Following certain methodological advances in the field of software studies, this dissertation uses new kinds of materials as historical sources: computer inputs, computer printouts, punched cards, numerical tabulations, and code.[21] Such sources permit analysis of how, for example, computers generated new definitions of disease and ways of handling medical information. This dissertation suggests that by attending to unfamiliar sources, from algorithms to computer inputs to numerical tabulations, historians can better understand how computers—and medicine's material grounding more broadly—have influenced the practice, experience, and work of medicine.

SITUATING COMPUTERISED DIAGNOSIS IN TIME

Before proceeding to the core discussion of computerised diagnostic systems and their connection to medical thought and practice, it is first important to demarcate where this historical study stands within broader histories of computing. Where and when in the history of computing did these diagnostic systems first surface?

[21] See, e.g., Matthew Fuller (ed.), *Software Studies: A Lexicon* (Cambridge, MA: MIT Press, 2008).

Different historical sources advocate for different starting points to the history of digital electronic computing. This historiographical discord speaks to the fact that the computer emerged not from the sharp mind and gifted hands of a single inventor, but from a loosely-defined period of rich intellectual activity. Within this diffuse period, however, it would not be unreasonable to pick out the year 1943. In the winter of that year, the first programmable digital electronic computer, the British Colossus, was deemed operational. Developed to assist British intelligence officers in breaking German codes, Colossus was as much a product of collaborative teams of scientists and technicians as it was a product of the broader exigencies of war.[22] Through the 1940s, a number of other computing devices joined the ranks of the British Colossus. Among the most prominent of such devices was the IBM Automatic Sequence Controlled Calculator, known more commonly as Harvard Mark I, which was used between 1944 and 1959 to compute tables for various branches of the American military.[23] The brainchild of Howard Aiken, a physicist and computing pioneer at Harvard, Mark I achieved significant cultural currency owing to its status as the first fully automatic electronic computer. Once programmed by inserting 'operation codes' punched onto a segment of paper tape, the machine was capable of running computations for days on end.[24] Such technical feats were enabled by an onerous paraphernalia of hardware: the machine stood fifty-one feet wide and two feet deep, weighed about five tons, relied on

[22] Indeed, intelligence obtained with the aid of Colossus informed the timing of the Normandy landings in 1944. Thomas Flowers, Alan Turing, William Tutte, and Max Newman are among the key individuals who worked on this early computer. B. Jack Copeland, 'Colossus: Its Origins and Originators', *IEEE Annals of the History of Computing* 26, no. 4 (2004): 38-45.
[23] Campbell-Kelly and Aspray, *Computer*, pp. 60-64.
[24] The speed at which the devices could perform these calculations was limited: three additions or subtractions took one second, a single multiplication took six seconds, and a trigonometric function was computed over the course of multiple minutes.

a five-horsepower electric motor, and was assembled from thousands of kilometres of wiring and nearly a million different parts.

From these earliest digital electronic systems, further improvements in computing technology did not proceed in a linear fashion—computers built just years after the British Colossus and Mark I were orders of magnitude faster and more complex than their progenitors. As one notable example, the Electronic Numerical Integrator and Computer (ENIAC), developed in 1946 by John Mauchly and J. Presper Eckert at the University of Pennsylvania, could perform five thousand operations in the time that Mark I could perform three.[25] While these capabilities gave the machine broad practical relevance, the ENIAC in fact emerged initially as a pragmatic solution to a narrow problem: to carry out time-consuming and tedious artillery firing table calculations as a way of predicting projectile trajectories under a set of external and environmental constraints. Humans at the Army's Ballistic Research Laboratory could only perform these laborious calculations at a certain pace, so the ENIAC's original and primary purpose was to relieve this human-resource bottleneck. Unlike the British Colossus or Mark I, the ENIAC was 'Turing-complete'—the device could, with sufficient memory, compute any conceivable calculation. Thus, Colossus, Mark I, and the ENIAC—as the first programmable digital electronic computer, the first automatic of these computers, and the first Turing-complete computer—embody three hallmarks in the history of digital electronic computing. With the collective appearance of these devices through the 1940s, machines began to supplant humans as the primary referents of the title 'computer'. As the mathematician and computer scientist George Stibitz

[25] For a detailed account of the events that culminated in the development of this computing device, see Harry Polachek, 'Before the ENIAC', *IEEE Annals of the History of Computing* 19, no. 2 (1997), pp. 25-30. Also see Campbell-Kelly and Aspray, *Computer*, pp. 76-82.

stated in a 1945 report for the National Defense Research Committee, 'By "computer" we shall mean a machine capable of carrying out automatically a succession of operations … and of storing the necessary intermediate results'. 'Human agents', Stibitz continued, 'will be referred to as "operators" to distinguish them from "computers" (machines)'.[26]

Although the historical significance of these devices was not fully recognised at the time, by no means did they fall upon deaf ears. Indeed, the ENIAC was featured on the front page of the *New York Times* in 1946, hailed as a tool that not only 'computes a mathematical problem 1,000 times faster than has ever been done' but also might even 'begin to rebuild scientific affairs on new foundations'.[27] Subsequent media coverage of the ENIAC, and its various counterparts, continued in this sensationalist vein, with new reports declaring, 'Computers Beat Brain: New Devices Said to be 100,000 Times Faster'.[28] Such popular responses fit within, and contributed to, a broader wave of technocratic optimism that swept through the United States beginning in the late 1940s and continuing into the early 1960s.[29]

This optimism was not wholly unfounded; these decades saw technologies of digital electronic computing take off. By the 1970s, computers only faintly resembled their precursors of two decades prior: once prohibitively expensive, computers became cheap enough to market to individuals as well as institutions; once unmanageably large,

---

[26] Quoted in Paul Ceruzzi, 'When Computers Were Human', *IEEE Annals of the History of Computing* 13, no. 3 (1991): 237-244, p. 240. In particular, 'computers' often referred to women. On the history of computing and gender, see, e.g., Janet Abbate, *Recoding Gender: Women's Changing Participation in Computing* (Cambridge, MA: MIT Press, 2017).
[27] T.R. Kennedy, Jr., 'Electronic Computer Flashes Answers, May Speed Engineering', *New York Times*, 15 February 1946, p. 1.
[28] 'Computers Beat Brain: New Devices Said to be 100,000 Times Faster', *New York Times*, 31 January 1947, p. 5.
[29] Edwards, *The Closed World*, p. 134.

computers of the 1970s were no longer the size of large rooms; and once limited by the constraints of vacuum-tube switches, the computer's processing power through these decades surged with the advent of transistors. The computer, however, did not progress from the mainframe to the personal computer in a sharp, punctuated fashion; rather, this progression unfolded gradually, through the development of various intermediate technologies and stages.[30] A cadre of changes across the computing landscape during the 1950s and 1960s—the development of transistors and integrated circuits, the introduction of time-sharing, the invention of the microprocessor—successively enhanced computers and computing technology along all three dimensions of cost, size, and speed. These technological innovations, in tandem with an emerging computer industry, fundamentally altered both the computer's material form in addition to its economic, professional, and social function. Unlike the first generation of digital electronic computers, which functioned exclusively in the service of mathematical computation, these subsequent iterations, many of which used integrated circuit assemblies, proved adaptable, their capabilities valuable across a range of domains. Clinical medicine stood as one of these domains.

Running parallel to this historical account of computer hardware and software is a story about technological advancements in machine learning and AI—both the technological advancements themselves and the broader societal perceptions of these advancements. Discussions about machine learning and AI have existed for as long as computers have. In fact, years before Turing's landmark 1950 paper on AI and the imitation game, cultural commentators were sizing the performance and capacities of computers up against human intelligence. Articles that assigned intelligence to early computers like Mark I and Colossus swarmed popular outlets of all kinds. In, for

[30] See chapter 8 of Campbell-Kelly and Apray, *Computer*, pp. 163-185.

example, a 1944 issue of *The American Weekly*, the magazine's science editor G.B. Lal invoked the idea of non-human forms of intelligence in deeming Mark I 'Harvard's Robot Super-Brain'.[31]

Despite this public rhetoric around AI, the field of computer science did not begin to make real concrete steps in AI research and development until the 1950s.[32] Among the first of these steps was the creation of specialised programming languages for AI. Between 1955 and 1956, three researchers affiliated with the RAND Corporation and Carnegie Mellon University—Allen Newell, Cliff Shaw, and Herbert Simon—designed the first of these AI languages, Information Processing Language (IPL), which introduced new features, such as list-processing and recursion, that allowed the language to function with an unprecedented degree of flexibility and to carry out general problem solving.[33] Later programming languages built upon the conceptual foundation laid by IPL. Most notably, in 1958, John McCarthy specified a new language distinguished by its melding of tree data structures and recursion.[34] McCarthy's List-processing language (Lisp) remained the hegemonic programming language in AI for decades after its development. The maturation of AI languages and

---

[31] G.B. Lal, 'Harvard's Robert Super-Brain', *The American Weekly*, 15 October 1944, p. 15. Public and academic perspectives on AI technologies has been an unsteady one, characterised by successive bouts of unbridled enthusiasm and intense disillusionment. For example, historians have identified periods during which major financing bodies, both private and governmental, significantly tapered funding toward AI research—a phenomenon these historians have termed an 'AI winter'. These stark periodisations, however, vastly oversimplify the research environment of AI over time.

[32] The history of AI also blurs together, and interacts, with the history of cybernetics. These histories are explored in Part I of this dissertation.

[33] Kurzweil, *The Age of Intelligent Machines,* pp. 198-199. The distinctive characteristics of IPL as a programming language—e.g., its linked-list formulation—were the products of Newell, Shaw, and Simon's efforts to computerise mathematical proof. Dick, 'Of Models and Machines'.

[34] Ibid., p. 202. Lisp remained the hegemonic programming language in AI for decades. John McCarthy, 'History of Lisp', in Richard L. Wexelblat (ed.), *History of Programming Languages I* (New York: Academic Press, 1978), 173-181.

technologies through the 1950s co-evolved with the professionalisation of AI as an academic and professional field: the first conference on topics related to AI was held in 1956, and the first research labs dedicated to AI research also cropped up within this time frame, including MIT's Artificial Intelligence Lab, founded in 1958 by Marvin Minsky and John McCarthy.[35] With these developments, the so-called age of intelligent machines began to take shape, and more sophisticated technologies like expert systems—computer programmes created to emulate the decision making processes of human experts—came into being.

In 1965, the Stanford computer scientist Edward Feigenbaum, together with the organic chemist Carl Djerassi and the Nobel laureate Joshua Lederberg, envisioned and constructed the first of these systems, DENDRAL, designed to assist organic chemists in identifying the structures of amino acids and new organic molecules (see Part III of this dissertation for more details).[36] As Feigenbaum described it, DENDRAL, and the rise of expert systems more broadly, marked a shift in the uses of computing technologies, with intricate and messy real-world concerns moving onto the AI agenda:

> During AI's first decade (1956-1966), the task environments in which AI scientists investigated their basic science issues were generally idealized "clean" task environments, such as propositional calculus theorem proving and puzzle solving. After the mid-1960s, a bolder and more applied inclination to choose complex real-world problems as task environments became evident. These efforts were both successful and exciting.[37]

The advent of early computerised diagnostic systems in fact preceded the historical demarcation set up by Feigenbaum. Yet medicine surely represented one of the

---

[35] Kurzweil, *The Age of Intelligent Machines*, p. 199, 202.
[36] Bruce Buchanan and Edward Feigenbaum, 'Dendral and Meta-Dendral: Their Applications Dimension', *Artificial Intelligence* 11 (1978): 5-24.
[37] Edward Feigenbaum and Bruce Buchanan, 'DENDRAL and Meta-DENDRAL: Roots of Knowledge Systems and Expert System Applications', *Artificial Intelligence* 59 (1993): 233-240, p. 233.

'complex real-world problems' to which he was referring. Few realms of medical practice are as confused, chaotic, and contingent as the realm of diagnosis. This dissertation follows the computer, an often rigid and rule-based device, into this realm.

THE PAGES THAT FOLLOW

Countless dissertations could have been written on the history of computerising medical diagnosis. The goal of this dissertation is less to offer a comprehensive reconstruction of history than to unpack important dynamics, tensions, and themes raised by the introduction of computers into diagnosis during the twentieth century. This dissertation thus opts for depth over breadth and adopts a case-based approach to historical analysis.

The dissertation proceeds in three sections, each comprised of two chapters and each organized around a particular case in the history of computerising diagnosis. Part 1, 'Paper Infrastructures', traces the creation of the Medical Data Screen (MDS), a computerised method of generating differential diagnoses. Though a novel digital method, the MDS in fact had its origins in a paper-based analogue technology: the Cornell Medical Index (CMI), a 195-item medical history questionnaire. Chapter 1 considers the creation and dissemination of the CMI: its motiving aims, its surprising uses, and its material grounding.

Over time Brodman's project of standardising medical history taking evolved into a project of computerising diagnosis. Through his work on the paper-based CMI, Brodman ultimately created the computer-based MDS. Chapter 2 explores the process of converting an analogue paper technology into a digital computerised one. The computerised MDS evolved from, and rested upon, paper data practices associated with questionnaires and surveys. Both tools were motivated by similar hopes. In an era of

overstretched and underresourced physicians, Brodman and his colleagues believed that these technologies would not only extend the reach of medical personnel to rural and underserved areas but also relinquish physicians from the monotonous aspects of medical practice, giving them more time to focus on the 'human' aspects of care. The CMI and MDS also stirred similar debates and tensions. The introduction of decision support technologies into medical history taking and diagnosis fuelled concerns around professional deskilling, interpersonal dehumanisation, and substandard care. Part 1 of this dissertation thus foregrounds the important continuities between computer-based digital practices in medicine and paper-based analogue ones.

Part 2, 'Recoding Disease', examines the role of the computer—and material culture more generally—as an agent of disease definition. The section explores this theme through the lens of a thirty-five year effort to computerise the diagnosis of haematologic diseases. In 1957, a group of researchers based at the Rockefeller Institute gathered at the Radio Corporation of America (RCA) facility in Camden, New Jersey to witness the RCA's Bizmac computer render, from inputs of patient data, diagnoses of haematologic diseases. Chapter 3 follows the development and evolution of this haematology programme. It shows how physicians and engineers, as they encountered novel digital media, began to represent medical information and define diseases in new ways. In exploring the computer's role as an agent of disease definition, this chapter suggests that the process of disease (re)definition is not just a social or intellectual process, as many historians have noted. It is a material process as well, shaped profoundly by the material media with which physicians work. Definitions of disease have a body—they are incarnate in stuff.

As interdisciplinary teams worked to automate medical diagnosis, they repeatedly bumped up against fundamental epistemological questions. How do doctors

think? What is the nature of human reasoning? Are human modes of reasoning amenable to computerisation? If they are, *should* computers aim to mimic them? In what ways do physicians come to define and classify disease categories? Are these categories even necessary? While following the haematology programme into its final decades, Chapter 4 explores how computers both generated and incorporated ideas about the medical mind and its inner workings.

Part 3, 'Black Box Medicine', uses the case of MYCIN to examine key moral tensions attendant to the computer's introduction into medical diagnosis. In the early 1970s, the young MD-PhD student Edward H. Shortliffe started work on his dissertation within Stanford University's Department of Computer Science. His PhD thesis, completed in 1975, involved the development of a computer programme called MYCIN, an expert system that provided physicians with advice concerning the diagnosis and treatment of bacterial infections. MYCIN immediately attracted wide attention: it matured into a much larger research enterprise at Stanford and spawned larger conversations about whether and how to apply computers to the problem of medical diagnosis. Chapter 5 recounts this early effort to computerise medical diagnosis and decision making. It demonstrates that questions about trust and transparency animated the development of, and responses to, the MYCIN project. How could a physician using the system know that its advice was accurate and trustworthy? What did the developers of MYCIN do to make the system's reasoning comprehensible to its human users? The MYCIN team largely pursued technological solutions to the problems of trust and transparency, creating novel technologies, programmes, and interfaces that purportedly allowed MYCIN to 'explain itself'.

As Chapter 6 argues, these issues around trust and transparency were entangled with concerns over authority, identity, and responsibility. To cede important treatment

and diagnostic decisions to machines—to allow computers to carry out a fundamental task of medicine—raised questions about what it meant to be a physician as well as questions about the physician's responsibility and his ownership over medical decisions. Despite all that the MYCIN did to render their system's reasoning transparent, physicians worried about lingering opacities of a computerised system—about the biases that, as one actor put it, may still be 'hidden in the code'. The technological fixes to issues of transparency thus fell short: they could not cleanly resolve the larger professional, regulatory, moral, and epistemological challenges raised by computerised medicine.

This is a dissertation about computerised medicine. Yet understanding the origins of computers in medicine requires a return to, and analysis of, period dominated by pencil and paper.

PART I

Paper Infrastructures

- CHAPTER I -

Indexing the World

Patients who came to the outpatient departments of the New York Hospital through the mid-twentieth century were accustomed to a familiar routine. They would check in at the front desk, be called in from the waiting room, and be seen by an attending physician, perhaps a medical resident. Some walked away, prescription in hand, while others awaited further testing. But patients who happened to come within an eighteen-month window—between July 1948 and December 1949—encountered a new and unusual step. These patients were asked to complete a four-page questionnaire. Emblazoned across the top of the questionnaire form were the words 'Cornell Medical Index'. Each patient spent roughly fifteen minutes answering the index's nearly 200 questions, which were designed to record the patient's complete medical history. For patients today, having to complete a health questionnaire seems par for the course. To patients in the late 1940s, however, this step was unprecedented: the Cornell Medical Index (CMI) was among the first health questionnaires of its kind.

The driving force behind the CMI was Keeve Brodman, a psychiatrist and professor at Cornell University Medical College (Figure 1—all figures are included after the main text of the dissertation). Together with a handful of colleagues, Brodman created the index as a method of improving medical history taking through standardisation. His early interest in standardisation soon evolved into an interest in computerisation. Beginning in the mid-1950s, Brodman and his collaborators began using data collected from the index to build a mathematical model that could diagnose a range of common diseases from a patient's questionnaire responses. Brodman's team then worked to convert this mathematical model into a programme for use on a digital

computer. The method that resulted from this research—the final version of which came to be known as the Medical Data Screen (MDS)—represented one of the earliest attempts to computerise medical diagnosis and decision making.

Brodman's work was influential in its time. The CMI, among the earliest and most prominent medical history questionnaires, reached hundreds of thousands of physicians and patients around the world, and the MDS represented one of the first attempts to computerise medical diagnosis. His work was published in leading medical journals and covered by the national press. Despite this, Brodman and his initiatives have received little historical attention.[38] In filling this gap, this chapter details the creation of the Cornell Medical Index and the following chapter unpacks how this project that started out as a paper-based method of medical history taking eventually morphed into a computer-based method of making differential diagnoses. Together, the chapters highlight the ways in which this notable effort to computerise diagnosis evolved from, and rested upon, earlier paper data practices associated with surveys and questionnaires. Any complete genealogy of biomedical computing and computerised medicine must take seriously the continuities between computer-based digital practices and paper-based analogue ones.

BRODMAN'S BEGINNINGS

New York City was all young Keeve Brodman knew. Born and raised in the city, Brodman attended the College of the City of New York, earned his medical degree from Cornell University Medical College, and joined Bellevue Hospital's Department

[38] Howard Bleich, himself a pioneer in biomedical computing, has written a short, encyclopedic overview of Brodman and the CMI. Howard Bleich, 'Keeve Brodman and the Cornell Medical Index', *MD Computing* 13, no. 2 (1996): 119-20, 122, 124. Brodman is also mentioned briefly in other histories of technology and biomedical computing. See Reiser, *Medicine and the Reign of Technology*, p. 215.

of Medicine, where he specialised in psychiatry. Beginning in 1934, he juggled two clinical appointments: one as an Associate Attending Physician at Bellevue, and the other as a physician to outpatients at the New York Hospital. To train at Cornell Medical College and its affiliated hospitals in the 1930s was to train within a vibrant medical ecology at the height of its transformation; during this period, Cornell Medical College not only grew significantly in its own right but also came to be embedded within a much larger network of New York medical facilities. In 1927, Cornell University and the New York Hospital, which had been formally associated since 1912, agreed upon the establishment of a joint medical centre. The resulting New York Hospital-Cornell Medical Center first opened its doors on 1 September 1932. The 27-story medical complex, donning a regal alabaster façade, towered over the surrounding Upper East Side neighbourhoods (Figure 2). Behind its impressive physical exterior, Cornell supported clinical and research activities that drew international recognition.

It would take the exigencies of war to unsettle Brodman's roots there. In late 1940, as World War II intensified in Western Europe, Brodman wrote to his local Army Corps area indicating his 'desire to co-operate with the armed forces in this national emergency'.[39] Less than two years later, Brodman received a notice from Adjutant General's Office of the War Department, certifying his military commission as a Major in the Army Medical Corps.[40] Soon thereafter, on 18 August 1942, Brodman reported for duty at the Station Hospital in Camp Lee, a notable Army post in eastern Virginia. There, he started off as the Chief of the Hospital's Cardiology Section but

[39] Keeve Brodman to Headquarters, Second Corp Area, 90 Church Street, NY, 4 Oct. 1940, Keeve Brodman, MD (1906-1979) Papers, Medical Center Archives of New York-Presbyterian/Weill Cornell, New York, box 7, folder 6 (hereafter cited in the form KBP/7/6).
[40] The Adjutant General's Office, War Department to Brodman, 3 Jul. 1942, KBP/7/1; Henry Cave to Brodman, 8 Jul. 1942, KBP/7/6.

quickly climbed the military ladder. In March 1943, Brodman was assigned the temporary position of Acting Chief of the Medical Service, and in June he took on a longer term appointment as the Medical Service's Assistant Chief.

Brodman made his role in the military continuous with his civilian medical role; he treated the Medical Corps as a new arena to which he could apply his knowledge of psychiatry. Early in his military tenure, Brodman cultivated an interest in psychosomatic medicine, a field that was just beginning to make inroads into the medical mainstream. This interest in psychosomatic medicine temporarily pulled Brodman away from his post in Virginia: in April and May of 1943 he attended a course in psychosomatic medicine with Dr. Helen Flanders Dunbar at Columbia University's College of Physicians and Surgeons. As someone with training in neuropsychiatry, Brodman understood the significance of working with and learning from Dunbar. By the early 1940s, Dunbar has established herself as a leading researcher and clinician in the field of psychosomatic medicine, adopting what one historian has called a 'holistic approach' to psychosomatic problems.[41] As the director of one of the first psychosomatic research programs, Dunbar collected holistic data on more than 1,600 patients at Columbia-Presbyterian Hospital in New York—their health, their family backgrounds, their social and economic living conditions. On the basis of this data she famously argued for the connection between psychosomatic disorders and certain 'personality constellations'. Having founded both the American Psychosomatic Association and its journal *Psychosomatic Medicine*, she was among the most widely recognised players in the field.

---

[41] R. C. Powell, 'Helen Flanders Dunbar (1902-1959) and a Holistic Approach to Psychosomatic Problems. I: The Rise and Fall of a Medical Philosophy', *Psychiatric Quarterly* 49, no. 2 (1977): 133-52; R. C. Powell, 'Helen Flanders Dunbar (1902-1959) and a Holistic Approach to Psychosomatic Problems. II: The Role of Dunbar's Nonmedical Background', *Psychiatric Quarterly* 50, no. 2 (1978): 144-157.

The Army granted Brodman six weeks of 'Detached Service' to go to New York and attend the course given by Dunbar.[42] During his stint in New York, Brodman interacted with Dunbar on a fairly personal level. In a letter he wrote to Dunbar after returning to Virginia, Brodman expressed his gratitude 'for the many cocktails' she gave him.[43] But Brodman also left New York feeling that these interactions landed somewhat short of what he had hoped for; he lamented in the same letter, 'I regret that the months of April and May were such busy months for you since that deprived me of one of the most important things I came to New York for—to hear your opinions on psychosomatic medicine'. Any disappointment over his limited engagements with Dunbar did little to curb his zeal for the nascent field. As early as May 1943, the same month that he returned from New York, Brodman began laying the groundwork for the establishment of a Psychotherapeutic Clinic in the Cardiac Section of Station Hospital.[44] The Clinic offered group psychotherapy to patients whose organic symptoms were determined to be psychogenic; Brodman estimated that such patients comprised 93 percent of all patients seen in the Hospital's Cardiac Section.[45] Through formal talks, Brodman also delivered the tools and tenets of psychosomatic medicine to his fellow medical officers.[46]

At Camp Lee, Brodman's psychosomatic programmes and reforms found a receptive audience. In a July 1943 communication to Brodman's commanding officer,

---

[42] Brodman to Arnold Albright, 29 Mar 1943, KBP/7/2; Brodman to Patrick J. Madigan, 11 May 1943, KBP/7/6.
[43] Brodman to Flanders Dunbar, 1 Jun. 1943, KBP/7/2.
[44] Brodman to Commanding Officer, Station Hospital, 'Group Psychotherapeutic Clinics for Patients in the Cardiac Section, Station Hospital, Virginia', 29 May 1943, KBP/7/2.
[45] Brodman to Commanding Officer, Station Hospital, 'Group Psychotherapeutic Clinics for Patients in the Cardiac Section, Station Hospital, Virginia', 29 May 1943, KBP/7/2.
[46] Memorandum, M. Chase to All Medical Officers, Station Hospital, Camp Lee, n.d. (circa June 1943), KBP/7/2.

Lieutenant Colonel M. Chase commended Brodman's service at Lee's Station Hospital: 'He … effectively assisted in the reorganizing of the Medical Service, establishing a regular medical clinical and consultation service, and establishing a psychosomatic clinic, which helped cut down the patient load at the hospital considerably'.[47] Another of Brodman's peers at Camp Lee wrote a month later, 'Keeve, I can never thank you enough for the splendid work you did for us here at Lee'.[48] Despite Brodman's successes locally, his efforts to introduce psychosomatic medicine more widely through the Medical Corps failed. Upon returning to Camp Lee from Detached Service in New York, Brodman began drafting a report to be submitted to the Surgeon General's Office (SGO) that proposed a psychosomatic teaching and consultation programme for the entire Medical Corps. From the outset Brodman was sober about the chances that the SGO would accept his recommendations. When, in July 1943, Brodman circulated a draft of the report to one of his superiors, he confided, 'I'm not too sanguine about the possibilities of it's [sic] acceptance; I console myself with the thought that although a lot of work goes into the preparation of such a report, one learns a lot'.[49]

The core of Brodman's proposal was a consult program, whereby two medical officers versed in psychosomatic medicine—one psychiatrist and one internist—would make rounds to military hospitals for in-depth and repeated consultations. Supplementing the consultation programme would be a teaching programme to 'acquaint the medical officer with the fundamental principles of psychosomatic medicine and morale that apply to military practice'.[50] In detailing the teaching program, Brodman touted a humanistic ideal of medical practice, one that highlighted

[47] M. Chase to Commanding Officer, Station Hospital, 8 Jul. 1943, KBP/7/2.
[48] Anon. to Brodman, 9 Aug. 1943, KBP/7/2.
[49] Brodman to Patrick J. Madigan, 11 May 1943, KBP/7/6.
[50] Keeve Brodman, 'A Program for Psychosomatic Medicine in the Army', unpublished report, p. 10, KBP/7/6.

the therapeutic importance of simple human interaction over sophisticated technological fixes. 'The fact will be stressed that "therapeutic procedure" often consists of a kindly word, a hopeful attitude, and a sympathetic listening, and is not limited to complex and delicate procedures which require long training and experience. These simple and often homely procedures combined with understanding are the sine qua non of all psychotherapy'.[51] Somewhat counterintuitively, Brodman geared the programme more toward medical and surgical officers than toward neuropsychiatric officers. Because medical and surgical officers were the first point of contact for patients with somatic complaints, Brodman felt that psychosomatic knowledge (about the ways in which somatic ailments cause and are caused by emotional phenomena) would do the most good in the hands of 'somaticists'.

Brodman defined the scope of psychosomatic medicine liberally. In his view, because all illness has psychological and emotional dimensions, psychosomatic medicine has something to offer all patients. Brodman explained in the report, 'Every soldier being cared for by the Medical Corps, no matter what the nature of his disease or injury, has a psychological factor associated with his illness. This factor may be an important etilogical [*sic*] agent, or may simply represent a reaction to the disease. It may be mild and relatively inconsequential, or of such severity that it plays a major role in the manifestations of the disease'.[52] By attending to the overlooked psychological and emotional facets of disease, Brodman argued, the Army would address the medical needs of soldiers more effectively and deploy its own medical resources more efficiently. Officials in the higher levels of the American military apparently did not agree with Brodman's assessment. On 17 July 1943, Brodman received confirmation of

[51] Brodman, 'A Program for Psychosomatic Medicine in the Army', p. 12.
[52] Ibid., p. 1.

his early fears: the SGO rejected the proposal, noting that '[t]here are not as many cases in the service that need psychosomatic therapy as this article would infer' and that, where such cases did exist, they would 'more properly come under the [existing] neuropsychiatric service'.[53]

When Brodman received this notice from the SGO he was in the midst of coming to terms with a more personal setback. Beginning in May 1943, Brodman began experiencing a series of peculiar symptoms. At first, they were slight: a feeling of tenseness in the legs, a sense of instability when climbing and descending stairs. But as the heat and humidity of a Virginian summer settled in over Camp Lee, his symptoms became harder to dismiss. In July, following an hour's worth of vigorous walking, Brodman became unusually fatigued and started noticing a loss of muscle control upon climbing some stairs. His legs grew progressively tired. By November 1943, Brodman was admitted to the Station Hospital of a neighbouring Army installation, Camp Pickett, where he was diagnosed with multiple sclerosis, a disease of the central nervous system (CNS) that variably affects a range of physical and mental functions. By early 1944, Brodman had been reclassified to 'limited duty' at the Walter Reed General Hospital in Washington, D.C. Reclassification, Brodman explained in one letter, 'is due to the recent development of multiple sclerosis, which for the time being causes considerable difficulty in walking even across the room'.[54]

Once reclassified, Brodman lobbied the SGO to be transferred back to New York—to an assignment with 'the Cornell Project,' a team of Cornell-based researchers working with the National Research Council for the study of psychoneurosis in the

[53] Surgeon General to Commanding General, Third Service Command, 17 Jul. 1943, KBP/7/2.
[54] Brodman to Surgeon General (War Department), 'Request for Assignment', 28 Jan. 1944, KBP/7/2.

Army.[55] Such an arrangement would fulfil Brodman's wishes both to remain in the Army and to pursue research of personal interest. On 22 January 1944, Harold Wolff, a Cornell neurologist and pioneer in the study of psychosomatic illness, wrote to Karl Menninger of the SGO in support of Brodman, detailing Cornell University Medical College's work on 'methods for aiding in the screening' of military personnel and 'request[ing] that Major Brodman be attached to this project' as a kind of bridge between civilian researchers and the U.S. Army.[56] In particular, Wolff described the group's role in devising 'the so-called Cornell Selectee Study which is … a brief neuropsychiatric and psychosomatic history having certain quantitative characteristics'. According to Wolff, the method allowed for 'quick and simple detection' of neuropsychiatric and psychosomatic problems that general medical officers often neglected.

But before any arrangement could be made for Brodman to join the Cornell Project in a limited duty capacity, administrative forces set Brodman on a different path. Difficulty in finding work that was suitable for limited duty active service and Brodman's disability prompted his summoning before the Army Retiring Board on 24 March 1944.[57] Before the Board, a medical examiner said of Brodman, 'He certainly is not qualitied for full duty, nor is he qualified for routine limited duty'.[58] It was the decision of the Board to put Brodman on terminal leave; about one month later, on 28 April, Brodman formally retired, transitioning to inactive status.[59] Although this was not the path Brodman had envisioned for himself, it ultimately brought him to his

[55] Ibid.
[56] Harold Wolff to Karl Menninger, 22 Jan. 1944, KBP/7/2.
[57] C. D. Bowen to Brodman, 18 Mar. 1944, KBP/7/2.
[58] Malin Craig, Proceeding of an Army Retiring Board in the Case of Major Keeve Brodman, Medical Corps, Army of the United States, Washington D. C., 24 March 1944, p. 3, KBP/7/2.
[59] T. W. Burke, Certificate, 29 March 1944, KBP/7/2.

desired destination: he returned to Cornell and began work on the National Research Council project studying psychoneurosis in the military.

THE CORNELL MEDICAL INDEX

Back at Cornell, Brodman nursed a number of different academic projects concurrently. One of these projects involved the study of psychiatric health and human relations in industry, which culminated in the 1947 book *Men at Work: The Supervisor and His People*.[60] Although Brodman's research on industrial health proved influential, reaching officials working at high levels in industry, his work on the 'Cornell Indices'—the Cornell Selectee Index and the Cornell Service Index—exerted a greater influence over the direction of his later career.[61] The Selectee Index dovetailed with Brodman's own interests and expertise. In line with Brodman's initiatives at Camp Lee, the Selectee Index aimed to furnish a simple, rapid, and accurate means of identifying psychoneuroses and psychosomatic disturbances in troops.[62] The Service Index simply included additional questions tailored for persons who had already undergone military training. Brodman's contributions to these initial Indices inspired him to develop a similar instrument except with a broader scope—the Cornell Medical Index.

---

[60] Keeve Brodman, *Men at Work: The Supervisor and His People* (Chicago: Cloud, 1947).
[61] T. B. Campion to Brodman, n.d., KBP/3/10.
[62] Arthur Weider, Bela Mittelmann, David Wechsler, Harold Wolff, and Margaret Meixner, 'The Cornell Selectee Index: A Method for Quick Testing of Selectees for the Armed Forces', *Journal of the American Medical Association (JAMA)* 124 (1944): 224–228; Arthur Weider, Keeve Brodman, Bela Mittelmann, David Wechsler, and Harold Wolff, 'Cornell Service Index: A Method for Quickly Assaying Personality and Psychosomatic Disturbances in Men in the Armed Forces', *War Medicine* 7 (1945): 209–213; Arthur Weider, Keeve Brodman, Bela Mittelmann, David Wechsler, and Harold Wolff, 'The Cornell Index: A Method for Quickly Assaying Personality and Psychosomatic Disturbances, to Be Used as an Adjunct to Interview', *Psychosomatic Medicine* 8 (1946): 411–413; and Bela Mittelmann and Keeve Brodman, 'The Cornell Indices and the Cornell Word Form: 1. Construction and Standardization', *Annals of the New York Academy of Sciences* 46 (1946): 573–578.

Although the CMI and its immediate ancestors were novel in their particulars, the larger impulse to survey and quantify populations was not. Tools of aggregation have long played important roles in American society—at least since the first U.S. census of 1790.[63] The 'quantifying spirit' took a firmer grip on American life through the late nineteenth and early twentieth centuries, as life insurance companies developed elaborate methods for counting, sorting, and measuring people—and for predicting their futures.[64] The logics and practices that animated corporate culture soon flourished in popular culture as well. Beginning in the interwar period, the proliferation of modern surveys—from the Gallup poll to the Kinsey Reports—generated not just large amounts of data on large segments of the population, but also new subjectivities.[65] Within medicine, the second half of the twentieth century saw the formation of new and significant links between the fields of medicine and statistics and, as a result of these links, the rising status of epidemiological studies and randomised controlled trials.[66] All the while, the creation and application of standardised tests and questionnaires, particularly within the fields of psychiatry and psychology, followed an upward trajectory through the twentieth century.

---

[63] Margo Anderson, *The American Census: A Social History* (New Haven, CT: Yale University Press, 1988).
[64] Dan Bouk, *How Our Days Became Numbered: Risk and the Rise of the Statistical Individual* (Chicago: University of Chicago Press, 2015). On the interrelations between medicine and life insurance, see Audrey Davis, 'Life Insurance and the Physical Examination: A Chapter in the Rise of American Medical Technology', *Bulletin of the History of Medicine* 55 (1981): 392-406. On the 'quantifying spirit', see Tore Frängsmyr, J. L. Heilbron, and Robin Rider, eds. *The Quantifying Spirit in the 18th Century* (Berkeley: University California Press, 1990).
[65] Sarah Igo, *The Averaged American: Surveys, Citizens, and the Making of a Mass Public* (Cambridge: Harvard University Press, 2007).
[66] Harry Marks, *The Progress of Experiment: Science and Therapeutic Reform in the United States, 1900-1990* (Cambridge: Cambridge University Press, 1997). See also Robert Aronowitz, *Risky Medicine: Our Quest to Cure Fear and Uncertainty* (Chicago and London: University Chicago Press, 2015), 69-94.

It was against this backdrop that Brodman and a number of colleagues (Albert J. Erdmann, Jr. and Irving Lorge chief among them) began work on the CMI in 1947 (Figure 3). The CMI differed from the Selectee Index in two primary ways. First, whereas the Selectee Index was designed for use with either members of or applicants to the Armed Forces, the CMI was written to apply to all patients. Second, the CMI aimed to generate a portrait of not just a patient's neuropsychiatric wellbeing but his 'total medical problem' as well.[67] In short, the CMI widened the Selectee Index's diagnostic parameters to encompass additional medical problems and patient populations.

Arriving at the right questions to include on the questionnaire was a two-year process of iteration and experimentation. The CMI's developers first tried drafting questions as if they were speaking to the patient. To Brodman's surprise, this strategy resulted in patient responses that he found to be 'completely inaccurate and unreliable'.[68] For example, a first draft of one question seemed simple enough: 'Were you ever a patient in a mental hospital?' But after testing the question with a small sample of patients, Brodman noticed that an incredibly large proportion of them responded affirmatively; many patients had apparently overlooked the word 'mental'. Through multiple rounds of revision, Brodman came to the question's formulation as ultimately included on the CMI: 'Were you ever a patient in a *mental* hospital (for your nerves)?'

---

[67] Keeve Brodman, Albert Erdmann, Jr., Irving Lorge, Harold Wolff, and Todd Broadbent, 'The Cornell Medical Index: An Adjunct to Medical Interview', *JAMA* 140 (1949): 530–534, p. 530.

[68] Keeve Brodman, 'The Cornell Medical Index: A Health Questionnaire to Help the Doctor Help You', unpublished article for *Cosmopolitan Magazine*, n.d. (circa July 1949), pp. 2-3, KBP/3/10.

Brodman, Erdman, Irving and their collaborators unveiled the completed CMI to the wider medical community in a seven-part series of articles published between 1949 and 1953.[69] The first of these, published in the *Journal of the American Medical Association (JAMA)*, introduced the questionnaire as 'a quick and reliable method of obtaining important facts about a patient's medical history without expenditure of the physician's time.'[70] A physician reading any one of these articles would be left with a good sense of the index and how it worked. The index included 195 questions that clustered loosely together into four different categories: questions concerning bodily functions, questions concerning past illnesses, questions concerning family history, and questions concerning behaviours, feelings, or moods. On the questionnaire form itself, these questions were arranged into eighteen sections (A though R) that spanned four pages. The first sections (A through L) mapped onto different organ systems, whereas the later sections (M through R) concerned disturbances of mood and feeling. In the style of its questions (and in the organisation of these questions on the form), the index replicated a common technique in medical history taking known as a review of systems

[69] Brodman et al., 'An Adjunct to Medical Interview', pp. 530–534; Keeve Brodman, 'The Cornell Medical Index-Health Questionnaire: II. As A Diagnostic Instrument', *JAMA* 145 (1951): 152–157; Keeve Brodman, Albert Erdmann, Jr., Irving Lorge, and Charles Gershenson, 'The Cornell Medical Index-Health Questionnaire: III. The Evaluation of Emotional Disturbances'*, Journal of Clinical Psychology* 8 (1952): 119–124; Keeve Brodman, Albert Erdmann, Jr., Irving Lorge, Charles Gershenson, Harold Wolff, and Todd Broadbent, 'The Cornell Medical Index-Health Questionnaire: IV. The Recognition of Emotional Disturbances in a General Hospital', *Journal of Clinical Psychology* 8 (1952): 289–293; Albert Erdmann, Jr., Keeve Brodman, Irving Lorge, and Harold Wolff, 'The Cornell Medical Index-Health Questionnaire: V. Outpatient Admitting Department of a General Hospital', *JAMA* 149 (1952): 550–551; Keeve Brodman, Albert Erdmann, Jr., Irving Lorge, and Harold Wolff, 'The Cornell Medical Index-Health Questionnaire: VI. The Relation of Patients' Complaints to Age, Sex, Race, and Education'*, Journal of Gerontology* 8 (1952): 339–342; and Keeve Brodman, Albert Erdmann, Jr., Irving Lorge, Jerome Deutschberger, and Harold Wolff, 'The Cornell Medical Index-Health Questionnaire: VII. The Prediction of Psychosomatic and Psychiatric Disabilities in Army Training', *American Journal of Psychiatry* 111 (1954): 37–40.

[70] Brodman et al., 'An Adjunct to Medical Interview', p. 530.

(ROS), whereby a physician would systematically ask a patient about subjective symptoms associated with the various organ systems (often proceeding in a 'head-to-toe' manner).

Worth noting is the way in which the CMI's standardising spirit was born out in its very aesthetic form.[71] Each question was assigned an identifying number. To ask 'Do you notice ringing in your ear(s)?' was to ask question 013. A patient who reported 'frequently feel[ing] faint' had circled 'Yes' next to the number 093. Although questions were grouped together according organ systems, these groupings appeared not under the written name of their respective system but under a simple letter. A physician who wished to examine a patient's visual system needed to consult the responses to questions under the heading '**A**'. As these letters proceeded alphabetically, the Cornell Medical Index really was, at least visually, a kind of index—an index of, as Brodman described it, 'the total patient'. Aside from having to write out one's name, address, age, and occupation, completing the questionnaire required virtually no writing and certainly no description or elaboration. Reviewing the CMI was therefore less a matter of interpreting written text than discerning binary markings associated with letters and numbers. In a conventional medical encounter, a patient might vocalise to the doctor his problems of exhaustion and fatigue. But on the CMI he would simply mark as follows:

I

… Yes No 119

[71] For a rich engagement with the aesthetic dimensions of medicine, see John Harley Warner, 'The Aesthetic Grounding of Modern Medicine', *Bulletin of the History of Medicine* 88, no. 1 (2014): 1-47.

In its appearance, then, the index was impersonal, sanitised, anonymous, universal.

The CMI's formatting, however, was not simply a matter of aesthetics. Questions were designed such that only 'Yes' indicated the presence of a symptom that the physician should follow up on. Brodman hoped that CMI users would be able to glean a sense of the patient's 'total medical problem' just by scanning the questionnaire and noting the conspicuous pattern of 'Yes' answers. A concentration of 'Yes' answers within a given section would suggest a medical problem localised to a particular organ system. A smattering of 'Yes' answers throughout the entire form might instead point to a medical issue of a more diffuse nature, whether an emotional disturbance or multi-system organic illness. Compared to the question of interpretation, the question of implementation was cake. The second article in the series noted bluntly, 'The test is administered simply by handing the form to the patient and asking him to complete it'.[72]

The CMI's alphabetical, indexical, and aesthetic character situates it within a much longer lineage of 'paper technologies' that have been deployed to help collect, organise, store, sort, communicate, and transmit medical knowledge.[73] Physicians have collected and organised large numbers of patient medical histories (*observationes*) in an indexical fashion since at least the early modern period, as J. Andrew Mendelsohn and Volker Hess have described.[74] The physician Wilhelm Fabry von Hilden (1560-1634), for instance, amassed six hundred patient histories, drawn from several decades of

---

[72] Brodman, 'As A Diagnostic Instrument', p. 152.
[73] See J. Andrew Mendelson, 'The World on a Page: Making a General Observation in the Eighteenth Century', in *Histories of Scientific Observation*, ed. Lorraine Daston and Elizabeth Lunbeck (Chicago: University Chicago Press, 2011), 396–420; Volker Hess and J. Andrew Mendelsohn, 'Case and Series: Medical Knowledge and Paper Technology, 1600–1900', *History of Science* 48 (2010): 287–314; and Volker Hess, 'A Paper Machine of Clinical Research in the Early Twentieth Century', *Isis* 109 (2018):473–493.
[74] Hess and Mendelsohn, 'Case and Series', pp. 288-289.

medical practice, and published these histories in a series of six volumes beginning in 1606. These volumes were synthesised into a single edition, the *Opera omnia*, which was published after Fabry's death in 1646. Because few readers could be bothered to read the *Opera omnia*, or even its predecessors, in all their length and detail, the books were alphabetically indexed, directing users to the pages that contained *observationes* of particular interest.

The indexing of texts like the *Opera omnia*, however, differed from the CMI in an important way. In this early modern case, indexing imposed order upon medical knowledge in a post-hoc fashion: patient histories that had been collected chronologically over time were retroactively arranged topically or alphabetically. The CMI, by contrast, was prescriptive: order was imposed at the very moment of knowledge collection. All patient histories collected by the CMI contained the same points of data, on the same topics, and in the same format. In its prescriptive and standardised character, the CMI belongs to another lineage of paper technologies that also stretches back hundreds of years. As early as the eighteenth century, hospitals began using 'registers', essentially lists of information on a piece of paper, to keep track of the patients entering the hospital, their identifying information, their occupation, the meals they had eaten, and the treatment they had received. Such registers arrived at the hospital largely by way of the broader practices of early modern state administration and accounting. Many physicians soon adopted this 'administrative model' of formatting for the purposes of medical history taking; they recorded patients' medical information not just in a wall of text but in tables, lists, and charts as well. As Hess and Mendelsohn have noted, the proliferation of prescriptive formatting of this kind rendered medical knowledge 'comparable and combinable'.[75]

[75] Hess and Mendelsohn, 'Case and Series', p. 293.

In much the same way, the CMI rendered thick, narratively rich medical histories more amenable to aggregation and manipulation, comparison and collation. Brodman and his colleagues could compare how patients with different backgrounds or diagnoses completed the questionnaire. What percentage of patients diagnosed with heart disease answered 'Yes' to item 042 ('Do you find it hard to breathe')? To what extent did patients of different genders, occupations, ages, or education levels answer the questionnaire in distinctive ways? Did 'neurotic' patients answer more questions on the index affirmatively? In collating CMI data to answer questions of this kind, researchers often conveyed their findings in tables and graphs. In an article on using the CMI in the evaluation of 'emotional disturbances', Brodman and his colleagues compared CMI responses varied among five different groups of patients: (1) patients admitted to medical and surgical out-patient departments of the New York Hospital, (2) the subset of these patients who were diagnosed as emotionally disturbed upon hospital examination, (3) psychiatric out-patients at the Brooklyn Veterans Administration, (4) random, ostensibly healthy, New York City men and women, and (5) newly employed New York Hospital workers. Comparing the response patterns among these groups, researchers were curious to know whether psychiatric and emotionally disturbed patients answered 'Yes' to a significantly greater number of index items. The Cornell team approached this question through the paper technique of tabulation: by creating a table with each column representing a different patient population and each row representing a sequentially higher percentage of 'Yes' responses, the researchers could visualize how these patient populations may have completed the CMI in ways that might be diagnostically useful. Visualising the data on paper helped researchers

conclude that thirty or more 'Yes' responses constituted 'presumptive evidence on the CMI of a patient having an emotional disturbance'.[76]

THE DILEMMA OF APPLICATION

The purpose of the series of CMI articles was not simply didactic; in addition to explaining the questionnaire and its clinical uses, these articles worked to build trust in both the specific questions posed by the CMI and the broader project of distilling medical history taking into a questionnaire. Brodman and his coauthors first went about establishing the CMI as a legitimate form of history taking by statistical means. In the 1949 article, the CMI developers examined a sample of 179 patients from the adult general medicine clinic of the New York Hospital Outpatient Department who had both taken the CMI and had their medical history taken by hospital physicians. Comparing the histories obtained by these two methods—one novel, the other routine—the researchers found that the CMI garnered a far more detailed picture of the patient than did the conventional history, recording patient data that the history had missed across every symptom domain. In the case of hypertension, for example, the CMI noted a past history of hypertension in nearly 25 percent of the patients sampled while the conventional history only picked up on a past history of hypertension in 10 percent of the sample.

That the CMI collected a larger quantity of patient data did not mean that this data was of higher, or even comparable, quality. To get at the accuracy of index's data, Brodman and his colleagues compared the occasions on which the patient data obtained by these methods overlapped. Finding that a high degree of correspondence between the methods (95 percent), the researchers concluded that patients respond to the CMI as

[76] Brodman et al., 'The Evaluation of Emotional Disturbances', p. 121.

honestly and accurately as they respond to existing techniques of history taking. Brodman and his colleagues expanded on these analyses in subsequent instalments to the article series. In a 1951 study, they reported that clinicians examining only CMI data made diagnoses that were often more accurate and complete than the diagnoses made through a hospital investigation of the patient's full medical chart.[77] The researchers were not advocating the use of the tool without further medical evaluation, but they presumably felt that evaluating the index as a standalone tool powerfully demonstrated its clinical utility.

To accompany and fill out this statistical evidence, Brodman marshalled narrative evidence, describing cases where the CMI outperformed the examining physician. In one reported case history, the CMI picked up a slate of notable symptoms—chest pains, palpitation, tachycardia, dyspnea, ankle edema, cramps in the legs, headaches, dizziness, faintness and muscular twitching—that had apparently evaded the examining physician's scrutiny. Brodman and his colleagues left open the question of '[w]hether the physician did not elicit and record these symptoms because of lack of time, or preknowledge of the diagnosis…'.[78] In another case history, a woman was admitted to the hospital 'complaining of a "stroke" of three months' duration, with impairment of speech and right-sided weakness'. Physicians ultimately identified three conditions—syphilitic aortitis, central nervous system syphilis, and cerebral vascular accident—as the root causes of the patient's affliction, but conventional practices and personnel did not arrive at this conclusion easily: it took five physicians, a laboratory test, and a referral to recognise a crucial piece of medical history (the woman's past history of syphilis) that the CMI had noted immediately.

[77] Brodman, 'As A Diagnostic Instrument'.
[78] Brodman et al., 'An Adjunct to Medical Interview', p. 532.

An official, fifteen-page CMI Manual condensed much of this information into an accessible digest and channelled it to the subset of physicians who were actually considering adopting the index. The manual performed three primary functions: informational, advisory, and political. Although the rise of scientific medicine through the late nineteenth and early twentieth century had embedded the ideal of standardisation into many areas of medical thought and practice, there was in the 1940s and 1950s nothing quite like the CMI—at least not on a large scale. As a vehicle for information, the manual conveyed to the uninitiated physician all major aspects of the CMI—its overarching structure, motivating logic, and constituent questions. In its advisory capacity, the manual instructed physicians on the ways in which the index as described could and should be used. The manual, for example, identified the single 'most effective method of using the CMI'.[79] This involved adopting the index as an adjunct to a more conventional oral interview and having the patient complete, and the physician review, the CMI prior to the interview. While the manual was prescriptive in the realm of implementation, it was notably deferential in the realm of interpretation: 'No rules can be given here on how to interpret symptoms', the manual stated. 'Interpretations must be based on the physician's training knowledge, experiences and insight'.[80]

The manual's recommendations and qualifications very likely reflected the honest judgement of its authors, but the document must also be considered in the context of a broader politics of clinical authority. Brodman understood that he had developed an instrument that intervened in an activity—medical history taking—long

[79] Keeve Brodman, Albert Erdmann, Jr., and Harold Wolff, *Manual: Cornell Medical Index-Health Questionnaire* (Cornell University Medical College, 1956), p. 5, KBP/1/3.
[80] Ibid., p. 6.

considered fundamental both to the practice of medicine and to physicians' sense of autonomy and identity. Statements about the physician's claim to interpretive authority can therefore be read as a way of navigating what some scholars have called the 'dilemma of application'.[81] In order to demonstrate a need for the questionnaire, the creators of the CMI had to persuade physicians of their own deficiencies in the area of medical history taking. But doing so also risked antagonising the very population they needed to accept the tool.

This dilemma also played out in the Cornell group's published articles. In illustrating the CMI's clinical value, Brodman and his collaborators found themselves straddling a fine line between promoting the index and impugning the physician. The developers of the index made measured claims about the CMI's clinical implementation, and they treaded lightly around topics concerning the capabilities and practices of the traditional doctor. The role of the CMI, the authors noted frequently, was that of an adjunct. By their account, the CMI proved most valuable when employed in combination with conventional modes of history taking: 'The most effective technic for use of the CMI as an adjunct to interview has been found to be for the patient to complete the CMI before he sees the physician and for the physician to review the CMI before he interviews the patient. The physician thereby gains a background of knowledge about the patient to orient him in interview'.[82] When the researchers inevitably touched on the topic of error in medicine, they privileged exogenous over endogenous explanations. The physician is busy, not inept, they implied: 'by calling attention to the patient's symptoms and significant items of past

[81] Michael Mulkay, Trevor Pinch, and Malcolm Ashmore, 'Colonizing the Mind: Dilemmas in the Application of Social Science', *Social Studies of Science* 17, no. 2 (1987): 231-256.
[82] Brodman et al., 'An Adjunct to Medical Interview', p. 533.

history, it [the index] assures that their investigation will not be overlooked *because the physician lacked time to elicit them*'.[83] But despite such assurances, implicit in evidence of the CMI's value and legitimacy were claims about the limitations of physicians and conventional medical methods. A corollary to the statement that the index outperformed the physician was the statement that the physician underperformed the index. Inevitably words like 'overlooked' and 'neglected' crept into their discussions of doctors.[84]

Unsurprisingly, then, some physicians objected to—or at least admonished against the hasty adoption of—the CMI. Critiques of the Index found their highest soapbox in a 1952 *British Medical Journal (BMJ)* editorial entitled 'The Patient Tells His Story'. The *BMJ* rehearsed critiques of the index that would become familiar: the editorial mentioned the tool's dehumanising potential ('It will be remarked that there are other criteria which establish with more certainty and with greater humanism the diagnosis of a neurosis') and its diagnostic blind spots ('it cannot indicate a symptomless early cancer, and the psychological disorders which it clearly reveals are superficial rather than profound, the neuroses rather than the psychoses').[85] Brodman interpreted the editorial as hostile. In his view the editorial's underlying contention was that 'the way to collect a history is not with a printed form but by the physician sitting down with the patient and questioning him'.[86] But in fact its tone was more agnostic than antagonistic; the *BMJ* editorial staff ultimately concluded that such critiques missed the index's point. 'The CMI is not intended to be the finished portrait: it is but the outline of a rough self-sketch'. The worm's-eye perspective of the hurried clinicians

[83] Ibid., p. 534. Emphasis added.
[84] Erdmann, Jr. et al., 'Outpatient Admitting Department of a General Hospital', p. 551.
[85] 'The Patient Tells His Story', *British Medical Journal (BMJ)* 4796 (1952):1246-1247, p. 1246.
[86] Brodman to Erich Meyerhoff, 20 Jan. 1976, KBP/5/1.

or medical specialist had something to gain from the bird's-eye perspective offered by the CMI:

> If it must be admitted that no general physician sitting in an out-patient clinic addresses 50 questions to the patient's psyche, it will also be agreed that no specialist makes 100 forays beyond his own territory, and no psychiatrist probes the soma 100 times. The CMI does all this with a minimal expenditure of the consultant's time. It tells the hospital doctor almost as much of the patient's complaints and the patient's personality as is known to a good family doctor.[87]

The editorial elicited letters from *BMJ* readers holding views across the spectrum. One physician, F.A. Pickworth, was so receptive to the idea of medical history questionnaires that their limited uptake at the time left him perplexed: 'In fact, it may be asked why such a method is not more widely used by doctors. Perhaps some delay is related to the sanctity of the doctor-patient relationship—about whose emotional aspects some of us are becoming more aware'.[88] To Graham Grant and R.A.N. Hitchens, resistance to the CMI's adoption would have made a lot more sense. In a joint letter, these physicians insisted that only a personal interview could effectively build trust between physician and patient, detect important unspoken symptoms, and allay the anxieties of worried patients. Particularly in emotionally charged contexts, they wrote, 'the questionary is found wanting'.[89]

Occasional squabbles aside, the index was more often ensnared in larger debates about the clinical value of questionnaires, surveys, and mass (multiphasic) screening techniques than targeted directly. One of such debates formed around the broader logic of questionnaires and indexes in medicine—a logic that one stroke patient

---

[87] 'The Patient Tells His Story', p. 1246.
[88] F. A. Pickworth, 'Correspondence: The Patient Tells His Story', *BMJ* (27 Dec. 1952), p. 1414.
[89] Graham Grant and R.A.N. Hitchens, 'Correspondence: The Patient Tells His Story', *BMJ* 4801 (10 Jan. 1953): 100-101, p. 100.

complained in the *New England Journal of Medicine (NEJM)* made the physician rigid, the patient unheard, and the relationship between the two 'adversarial'. 'A rigid doctor intensifies a rigid patient', this patient lamented.[90] Discussing a medical textbook, he continued, 'A certain logic is discernible here, all right, but it is an indexer's logic, and it does chop up the stroke patient … into rather small pieces'. Mass screening programs, whereby apparently healthy people are examined for disease, had their detractors too. That they yielded too many false negatives and false positives, that they employed faulty tests, and that they left no favourable dent on patient morbidity and mortality were all criticisms that emerged as these screening programmes were gaining traction in the 1950s.[91]

Other seeds of criticism were planted in the empirical literature. A 1954 study on medical questionnaires in the journal *Clinical Science* suggested that patients tend to exaggerate the presence of symptoms when those symptoms are mentioned by the questionnaire.[92] This study's findings were discussed widely, occasionally in the context of the CMI. As stated in the leading article for a 1954 issue of the *BMJ*: '[I]t is becoming increasingly popular in the U.S.A. for a patient to provide his medical history of filling in a questionary such as the Cornell Medical Index. … Galser and Whittow, in light of their findings, issue a warning that the questionaries themselves may produce symptoms'.[93] Brodman apparently felt these critiques deeply. In 1965, he confessed to one of his more regular interlocutors, 'For the past twenty years I too have devoted

---

[90] Eric Hodgins, 'Listen: The Patient', *New England Journal of Medicine (NEJM)* 274 (1966): 657-661.
[91] Multiphasic screening programmes were also met with praise. Editorial, 'Report on the Screening Clinic', *NEJM* 243 (1950): 275; Vlado A. Getting and Herbert L. Lombard, 'The Evaluation of Pilot Clinics—The Mass Screening or Health-Protection Program', *NEJM* 247 (1952): 460-465.
[92] E. M. Glaser and G. C. Whittow, 'Experimental Errors in Clinical Trials', *Clinical Science* 13, no. 2 (1954): 199-210.
[93] 'Questionaries in Clinical Trials', *BMJ* 4875 (12 Jun. 1954): 1366-1367.

myself to a controversial and generally unpopular effort and have endured the loneliness you now experience'.[94]

But counterbalancing the criticism and concern among some medical professionals was an interest in and enthusiasm for the index among others. Hundreds of physicians from around the world wrote to Brodman commending him for his work on the CMI. One physician noted how he was 'keen to use the Cornell Medical Index as an important instrument of inquiry' in his own practice of medicine.[95] Other early adopters remarked on the tool's utility, reporting to have 'obtained good help' from the CMI.[96] Coverage of the CMI in the lay press also carried a distinctly positive valence. In 1951, the *New York Herald Tribune* ran the story, 'Questionnaire Helps Doctors Diagnose Ills', which introduced readers to the index and basic data about its clinical effectiveness. Citing this data, the article concluded, 'Since this questionnaire proved so helpful in aiding physicians … the questionnaire is being introduced in Bellevue Hospital and in numerous private hospital practices'.[97] A 1953 article in the Spanish-language magazine *Vision* similarly hailed the index as a remedy to '[o]ne of the main headaches plaguing today's busy doctor's [sic]'.[98]

INDEXING THE WORLD

[94] Brodman to H. J. van Heyst, 23 Feb. 1965, KBP/3/13.
[95] J. J. Fleminger to Brodman, 20 Jan. 1965, KBP/3/13.
[96] Harold Braun to Brodman, 26 Sep. 1966, KBP/4/1; Henry Glah, Jr. to Brodman, 9 Oct. 1966, KBP/4/1.
[97] Harriet Jean Anderson, 'Questionnaire Helps Doctors Diagnose Ills; Cornell Medical's Health Form Proves Efficiency During Four-Year Test', *New York Herald Tribune* (New York City, NY), 9 Mar. 1951, p. 16.
[98] Margaret Boomer to Brodman, 14 Apr. 1953, KBP/6/2, The quote from the introduction to an article—enclosed in the letter—that was to appear in a 1 May 1953 issue of *Vision*.

Enthusiasm for the index was backed by its actual clinical use. The distribution and sale of the CMI experienced healthy growth in the decades after its publication. Through the 1950s and 1960s, Cornell Medical College sold and distributed around a quarter of million copies of the index annually. Thousands more spread through networks of physicians informally. By the 1960s, the CMI was recognized in medical journals as the 'most widely used' medical questionnaire to date.[99] In addition to its broadening influence and visibility as a tool of clinical practice, the index also enjoyed a career as a tool of medical research. From 1949 onwards, the number of published articles employing the index in some fashion trended upwards. As Brodman and his collaborators pitched it, further research with the CMI would enhance not just medical knowledge but the value of the tool itself. Since the standardised nature of the CMI allowed comparisons across different studies and patient populations, a growing bank of published data derived from the index would incentivise its adoption by other researchers; as the size of the CMI research network grew, so did the expected value of joining that network. Many also speculated that additional CMI research would enrich the index's *clinical* value. As the *BMJ* editorial observed, future research might shed light on yet-undiscovered connections between questionnaire responses and medical conditions.

National and linguistic boundaries hardly inhibited the index's expansion through communities of clinicians and researchers. The CMI gradually became international and multilingual. The first translation of the index—into Spanish—came the early 1950s, at the behest of officials at Cornell and the New York Hospital. The New York Hospital served a large population of Puerto Rican patients, many of whom

[99] J. Anderson, 'New Self-Administered Medical Questionary', *BMJ* 4 (1968): 636-638.

came to New York City in the wave of migration to the city following World War II. Brodman therefore commissioned a group of professors at the University of Puerto Rico to make the Spanish translation.[100] Before long, the Cornell Medical Index was not just the *Indice Medico de Cornell* (Spanish); it also became the *Index Médical de Cornell* (French); *Mục-lục Y-khoa Cornell* (Vietnamese), and *Cornell Tıbbi Endeksi* (Turkish). Japanese, German, Arabic, Dutch, Hebrew, Hindu, and Chinese translations also made the Index intelligible to patients speaking these languages (Figure 4). By and large, these subsequent translations emerged organically—not through the directive of Brodman or Cornell but through the initiative of physicians and researchers working in far-flung regions of world. One physician who lived and worked in the Lago Colony, a small multinational community located on the east end of the Caribbean island of Aruba, presented Brodman in 1955 with a translation of the index into the island's local language, Papiamento. The physician estimated that this Creole language—derived from Arawakan, English, Dutch, Spanish, French, Portuguese, and a number of African languages—was spoken by just 180,000 individuals across the Dutch Caribbean islands.

At its core, the CMI was an instrument of standardisation. No two physicians confronted with the same patient would take the exact same medical history. Indeed, such variation was, in the early 1950s, emerging as a topic of inquiry—and concern. In 1951, Archibald Cochrane, together with two other British researchers, documented wide variations in the medical histories taken from comparable groups of coalminers, even when the medical observers asked the miners the exact same questions.[101] The authors hypothesised that much of this variation resulted from the 'bias' of the medical

[100] Brodman to Erich Meyerhoff, 11 Feb. 1976, KBP/5/1.
[101] A. L. Cochrane, P. J. Chapman, and P. D. Oldham, '"Observers" Errors in Taking Medical Histories', *Lancet* (5 May 1951): 1007-1009.

observers coming into contact with the unstructured nature of the miners' answers. When asked, 'Have you a cough?' or 'Do you spit?', the participants might reply imprecisely, 'Only an ordinary cough' or 'Only after work'.[102] How a physician codes, or follows up on, these responses inevitably turned on factors—experience, knowledge, background, personality—unique to each individual clinician. Open-ended history taking, Cochrane and his co-authors implied, leaves the door open to inconsistency and error. The index, by contrast, imposed sharp boundaries on the possible form a medical history could take. It asked all patients the same questions, in the same order, and in the same fashion; the responses available to patients were binary: 'Yes' or 'No'. The CMI manual, moreover, provided clear instructions concerning the Index's clinical implementation, demarcating the 'intended' (correct) from the 'unintended' (incorrect) ways of using the questionnaire.

Despite references in print to a single, universal, and standard CMI, in practice there existed a multiplicity of CMIs. As the CMI pushed through geographic and linguistic barriers, the tool sometimes changed along the way, adapting to its various local contexts. Even within the United States, some physicians who integrated the CMI into their private practices tailored the questionnaire to the needs of their practices and their patients. One internist wrote to Brodman in 1966, 'I have used the old Cornell Medical Index beginning about 1951 and a personal modification since 1961. I find these extremely useful in my practice of Internal Medicine'.[103] But not all changes to the index were intentional: meaning was also lost—and gained—in translation. The creators of the index were very deliberate about the wording of the 195 items on the questionnaire. The slightest modification of word choice, Brodman discovered in the

[102] Ibid., p. 1008.
[103] Philip Ball to Brodman, 28 Oct. 1966, KBP/4/1.

process of creating the CMI, could significantly alter patients' responses. 'We found', Brodman wrote in 1949, 'that every word had to be chosen with the greatest care, and that every question had to be rigorously tested after it was put together'.[104]

The chief criterion for choosing between different formulations of a given question was clarity; the CMI's developers wanted to design a questionnaire that could be answered by people with 'no more than three or four years of elementary school education'.[105] For Brodman, making the index linguistically accessible meant making it conversational and idiomatic—incorporating common expressions and dressing many questions up in a colloquial style. For example, item 145 on the index read, 'Does your work fall to pieces when the boss or a supervisor is watching you?' Idioms like these vary across languages, and as a result maintaining the 'spirit' of such questions sometimes proved difficult.[106] Once translated into another language, the questionnaire was never equally befitting of all dialects within that language. The Cornell-sponsored Spanish translation, customised for speakers of Puerto Rican Spanish, was, to speakers of certain other dialects, intelligible but imperfect. A Venezuelan doctor raised these issues to Brodman in a 1962 letter: '"Eyeglasses" are called "espejuelos" in Cuba. In Venezuela we called them "lentes" which would be more widely understood. "Anteojos" used in Colombia is also more clear than "espejuelos"'.[107]

Dwarfing the challenge of translating across languages was the challenge of translating across cultures. Built into the CMI were certain theories of illness along with certain assumptions about the population's health behaviours, views on medical

---

[104] Keeve Brodman, 'The Cornell Medical Index: A Health Questionnaire to Help the Doctor Help You', unpublished article for *Cosmopolitan Magazine*, n.d. (circa July 1949), KBP/3/10.
[105] Brodman to F. Chicou, 29 Mar. 1971, KBP/5/1.
[106] Hirotugu Miyaki to Brodman, n.d. (circa March 1967), KBP/4/2; Matilde de Dilva to Brodman, 25 Mar. 1953, KBP/6/2.
[107] R. D. Verson to Brodman, 12 Jun. 1962, KBP/6/2.

disclosure, and distribution of disease. Researchers found that these assumptions did not hold in all contexts. In the early 1950s, the psychiatrist Ludwig Laufer used the CMI among people native to the Okinawan islands located to the south of Japan. From the outset, Laufer identified an apparent mismatch between the index and the local culture: whereas the CMI was oriented toward the individual, Okinawan culture (in his view) was oriented toward the group. According to Laufer, seldom did a local Okinawan make decisions alone; instead he contributes to the decisions made by a family unit. Laufer therefore administered the index not to individuals but to 'cooperative family group[s]' that discussed and responded collectively to each questionnaire item. Compounding issues with the CMI that were more structural in nature, individual items on the questionnaire were often incongruous with Okinawan customs and habits. In a culture with a different family and social structure, what was the meaning of a question like, 'Do you wish that you always had someone at your side to advise you'? Or, more materially: given the higher prevalence of endamebiasis and intestinal parasites, what was Laufer to make of questions about respondents' bowel movements?

About a decade later, an anthropologist-physician team from Harvard faced similar problems in applying the index to Zulu communities in Southern Africa. The team created a modified version of the CMI, (modestly) adapted to the linguistic and cultural particularities of the population studied. These slight modifications did little to rescue a number of questions from cultural opacity. According to the study's authors, it was a common assumption in Zulu communities that strangers were *umtagathi*, or sorcerers. As a result, an affirmative answer to the question, 'Do strange people or places make you afraid?', was indicative not of any paranoid or neurotic mentality but of a perfectly normal cultural attitude. Beyond problems with individual items, the

entire structure of the questionnaire came under strain. The index's organisation into organ systems, based on Western ideas about disease causation and the human body, came into conflict with other theories of illness and systems of belief. Zulu indigenous culture, the researchers reasoned, 'provides a competing and highly organized system of classifications of disease and concepts of disease etiology, so that the classification of CMI questions into organ-systems or by disease areas may seem, to them, to present symptoms in a meaningless or confusing context'.[108] It was not the conclusion of these researchers that the CMI had buckled under the weight of these cultural forces; rather, the team suggested that the index proved useful in unintended ways. They proposed that the index and culture had a bidirectional relationship: the CMI must be interpreted with adequate knowledge of the participating population's culture, but they came to believe that, when interpreted as such, the CMI 'yields clues' not just to the respondents' health status but to their 'cultural and social situation'.[109]

The CMI's national and international spread highlights the role of the instrument's users in driving its clinical expansion and shaping its clinical use. Brodman himself recognized the possibility that, once developed, the tool might mature in unexpected, even regrettable, ways. As he wrote in a draft of 1949 magazine article, 'An[d] now, we're wondering about the future. Will the Health Questionnaire be used properly? Will it be recognized for what it is, simply a large collection of medical data for the doctor to study and combine with other information? Or will it be used like a radio tube tester, and all by itself, in an attempt to make automatic diagnoses?'[110] But

[108] Norman A. Scotch and H. Jack Geiger, 'An Index of Symptom and Disease in Zulu Culture', *Human Organization* 22, no. 4 (1963-1964): 304-311, p. 311.
[109] Scotch and Geiger, 'An Index of Symptom and Disease in Zulu Culture', p. 311.
[110] Keeve Brodman, 'The Cornell Medical Index: A Health Quuestionnaire to Help the Doctor Help You', unpublished article for *Cosmopolitan Magazine*, n.d. (circa July 1949), pp. 2-3, KBP/3/10.

in discussing the role of the index's 'consumers', I do not mean to diminish from the work of its producers in promoting and legitimating the tool. Brodman and his collaborators had grand ambitions for the CMI, and they chased these ambitions via a variety of means. To publicise the index, these researchers reported on the tool in leading medical journals; to facilitate its clinical implementation, they created a user manual that addressed logistical issues surrounding its clinical use. This first wave of publications on the index relied on a single patient population: outpatients at the New York Hospital. This limited subject pool satisfied these early studies' primary aim of illuminating the index's basic value in more structured and conventional clinical settings. But by the early 1950s, Brodman and his colleagues had initiated a number of studies to demonstrate the CMI's utility in different, sometimes even non-clinical, environments.

THE CMI GOES TO ONEONTA, NY

One of these studies, the Oneonta Health Survey, was ambitious. The study, conducted in November 1951, surveyed 1,800 adult residents from 600 randomly selected families in Oneonta, a town in central New York State. The study was a logistically complex operation, and the Cornell researchers enlisted the Bureau of Applied Social Research at Columbia University to assist in coordinating and undertaking the study's fieldwork. The Bureau's staff instructed and supervised twelve fieldworkers, drawn from the local Oneonta population, who distributed copies of the CMI to each of the study's adult participants and gathered additional medical, social, personal, and economic information about the participants from the head of each household. The fieldworkers directed participants to return the completed index by mail

and personally revisited participants who, after two weeks, failed to do so. By the end of the study, 85 percent of the study had responded.

The American public was not always well acquainted with surveys, questionnaires, and population statistics, as Sarah Igo has shown.[111] Brodman and his collaborator Irving Lorge made no assumptions about the Oneonta public's familiarity with—or their warmness toward—either the CMI specifically or health questionnaires in general. Cornell researchers orchestrated an information campaign not only to raise awareness about the health survey but also to put the survey under a positive light. The campaign, as Harold Wolff related in a 1952 letter, began ten days prior to the fieldwork and recruited the cooperation of 'local physicians, Public Health Service personnel, civic leaders, school authorities, newspapers, and radio stations'.[112] The campaign's various organs worked together under a common slogan: 'To help the doctor help you'. The 'you', Brodman explained to Lorge in the month leading up to the study, 'is both singular, referring to the individual completing the form, and plural, the community'.[113] To convincingly claim that the study would help the local community, the study's architects ironically needed the local community's help. 'We were always going to accent that this is a community affair', Brodman wrote, 'that the civic organizations, schools, churches, physicians, health officers, newspaper, and community leaders were backing the project'. Brodman and his colleagues recognised that the campaign's message would be amplified if channelled through the voice of local officials, institutions, and community members: 'The publicity should mention the fact that the community can take pride in being selected for this survey. Statements

[111] Igo, *The Averaged American.*
[112] Harold Wolff to Paul B. Magnuson, 29 Feb. 1952, KBP/3/5.
[113] Brodman to Irving Lorge, 2 Oct. 1951, KBP/3/5.

to this effect from local leaders are probably the most effective method of making this point'.

The formal objective of the Oneonta Health Survey was, as stated in some July 1951 notes on the project, '[t]o determine the health status of the population of Oneonta N.Y. and the contiguous rural area'.[114] But behind this primary research objective lay a grander secondary objective. In effectively appraising the health status of Oneonta, the developers of the CMI hoped that the study would prove the CMI's suitability for larger scale, even national, health surveys. Researchers selected Oneonta over alternative locations largely for the reason that the city, with its mix of urban and rural areas, seemed nationally representative; success in surveying the health Oneonta, these researchers figured, would augur well for the index's successful application at a national level. In December 1951, just months after the survey had wrapped up, Harold Wolff saw an opportunity to introduce the index into discussions about national health initiatives. That month, President Harry Truman signed an Executive Order establishing the President's Commission on the Health Needs of the Nation. The President left the Commission with a broad directive: 'During this crucial period in our country's history it [the Commission] will make a critical study of our total health requirements, both immediate and long-range, and will recommend courses of action to meet these needs'.[115]

Part of estimating the nation's total health requirements involved estimating its total health status. How many Americans had heart disease? What percentage complained of vision problems? This is where Brodman, Wolff, and their associates

[114] Notes on Oneonta Project, 31 Jul. 1951, KBP/3/5.
[115] President's Commission on the Health Needs of the Nation, *Building America's Health: A Report to the President* (Washington, D.C.: Government Printing Office, 1952), p. ix.

thought that CMI could be of use. In January 1952, Wolff wrote to Commission member Joseph Hinsey in an effort to bring the questionnaire to the Commission's attention. Wolff singled out the group's work in Oneonta as being of particular interest to the Commission. He opened the letter, writing, 'As a member of the commission … you may be interested in our experiences with the Cornell Medical Index-Health Questionnaire, especially those in making a health survey of an upper New York State community'.[116] He closed on a similar note: 'Our experiences in this local [Oneonta] survey have convinced us that the Cornell Medical Index-Health Questionnaire may be useful in conducting a national health survey'. Wolff's suggestion indeed piqued the Commission's curiosity. Commission chair Paul B. Magnuson followed up with Wolff in a series of letters, all of which homed in on certain administrative aspects of collecting health data with the index. Magnuson was eager to know how many people were surveyed in Oneonta, how long it took to conduct the survey, how long it took to analyse the information collected, how much money conducting a health survey of this nature cost. '[T]hese questions', Magnuson explained, 'are important to us in determining whether or not this Commission should sponsor such an undertaking on a larger scale'.[117] The Commission ultimately decided against using the index, curbing hopes that the CMI might be adopted as a government-backed, nationwide survey. To be sure this was a disappointment, but not one that stopped Brodman from pushing the CMI in new directions.

[116] Harold Wolff to Joseph C. Hinsey, 14 Jan. 1952, KBP/3/5.
[117] Magnuson to Wolff, 18 Feb. 1952, KBP/3/5.

- CHAPTER 2 -

The Statistical Patient

In January 1955, J. Robert Oppenheimer, who was then the director of Princeton's Institute for Advanced Study (IAS), appeared on Edward R. Murrow's popular television series *See It Now* to discuss the IAS and its work. When Brodman watched the program, he noticed two things: first, Oppenheimer's commitment to 'interdisciplinary thought in science'; and second, Oppenheimer's personal connections to the world's most prominent scientists and mathematicians.[118] Brodman therefore wrote to Oppenheimer asking if he might recommend a mathematician 'with competence and interest' in a problem that confronted Brodman. His problem was this. The fact that the CMI (1) aided physicians (and even unskilled personnel) in making comprehensive diagnostic decisions, (2) quickly collected large amounts of data on large numbers of patients, and (3) collected this data in a binary form suggested to Brodman that it was possible 'to derive a mathematical model of the operation of making diagnoses from the data, as an analogue to human thinking'. Although Brodman believed that he had identified an auspicious problem area, he recognized his own limitations in navigating it solo: 'My own inadequacy as [a] mathematician has permitted the construction of only a crude model, but even this Simple Simon gives promise that further study may be fruitful'.[119]

For Brodman, the hope of developing a diagnostic model from CMI data was inextricably linked to the hope of shedding light on the nature of human reasoning. From the beginning, the idea that a mathematical model could both incorporate and

[118] Brodman to J. Robert Oppenheimer, 12 Jan. 1955, KBP/3/11.
[119] Ibid.

illuminate human cognitive processes consumed him; in fact, this was his primary interest. As he noted to Oppenheimer, 'Study of such a model may have not only value in pure science but may also give information about the mechanisms of higher brain functions such as memory, concept formation, and judgement'.[120] By way of response, Oppenheimer advised Brodman to contact Julian Bigelow, a cyberneticist, engineer, and computing pioneer based at the IAS.[121] Brodman did just that. In a letter to Bigelow, Brodman highlighted his team's finding that non-physicians—medical interns, nurses, and technicians—made diagnostic deductions from CMI data nearly as effectively as physicians did. Even a layman, Brodman elaborated, 'with not as much pertinent information as a physician, can usually make partial diagnoses'.[122] To Brodman this hinted that (contrary to the CMI manual's insistence otherwise) the secret ingredient to interpreting patient data was not to be found exclusively in the physician; perhaps a mathematical model, based on human reasoning generally, would suffice. Brodman posed this question to Bigelow: 'since a human being with this information in binary form can make correct diagnostic decisions, is it not possible to derive a mathematical model of the operation of making these decision'?

Brodman offered a rough sketch of how, by his estimation, humans reason in a clinical context: 'Some of the methods a person uses in reaching decisions for diagnoses are known…. For example, in determining which organ system is likely to be associated with a complaint on the questionnaire, the diagnostician unconsciously scans his memory. He estimates the frequency with which the complaint occurs with disorders in each organ system, and selects the one with which the complaint is most

[120] Brodman to Oppenheimer, 12 Jan. 1955, KBP/3/11.
[121] Katherine Russell to Brodman, 3 Feb. 1955, KBP/3/11.
[122] Brodman to Julian Bigelow, 10 Mar. 1955, KBP/3/11.

frequently associated'.[123] Bigelow, having spent much of his career thinking about the application of computing and mathematics to practical domains of this sort, was stimulated by these ideas. He agreed that the proposal warranted further serious study, but he gently implied that Brodman was not fully appreciating the complexity of human cognitive processes: 'This is not at all easy to do in a valid way, as the scarcity of genuine successes reported in the literature may indicate…'. Like Brodman, Bigelow believed that there existed a mechanism by which humans make decisions from observational data, but in his view this mechanism was 'of the most complicated and obscure type'.[124]

In Brodman's mind, the complexity of human cognitive functions depended on the scale of analysis. At lower, more basic levels of cognition, Brodman insisted on simplicity. Either a neuron fires or it doesn't: 'The unit mechanism is almost certainly simple, a yes discharge or a no discharge. The next level for some functions may still be simple neurones arranged in configurations permitting such operations as 'and,' 'or' and 'not.' Then I stop'.[125] Beyond this stopping point, Brodman conceded that speculating about the mechanism of higher-level thought processes was difficult, noting 'the immense and bewildering complexities of interrelated functions'. But this complexity spurred Brodman. He saw the proposed study as a possible entry 'into this vast unknown field of knowledge'.

Brodman's view of human thought processes was refracted through the lens of his own experience with multiple sclerosis. Once Brodman returned to New York from his military service, his condition worsened. From 1948 onwards, Brodman was effectively housebound; when he met with correspondents and collaborators like

[123] Ibid.
[124] Bigelow to Brodman, 13 Apr. 1955, KBP/3/11.
[125] Brodman to Bigelow, 26 Apr. 1955, KBP/3/11.

Bigelow, it was always at his home. In 1954, Brodman more fully elaborated on the nature and progression of his symptoms: 'For the past dozen years I have developed increasing disabilities due to multiple sclerosis. Their marked progression during the past six months has made it extremely difficult for me to keep up with my work or, for that matter, to carry out the daily routine of living'.[126]

But Brodman's disability defined as much as it distracted from his work: it was his experience with multiple sclerosis that first convinced him that human reasoning could be formalised, mathematised, and computerised. Brodman began deriving ideas about human cognition from his neurodegenerative disease as early as 1953. In a 1953 letter to Norbert Wiener, Brodman described his own experience with multiple sclerosis in terms that Wiener, as the 'father of cybernetics', would have understood well—Brodman revealed how central nervous system tracts degenerate 'as if small groups of wires in a communication system were cut'. 'This suggests', he continued, 'that a study of functions disordered by the disease may give useful leads to the understanding of central nervous system mechanisms. I am thinking not only about spinal and cerebellar functions but more particularly about the unknown mechanisms of such activities as memory, association, judgement, intellegence [sic], and behavior'.[127] Two years later, in a letter to Bigelow, Brodman invoked this same cybernetic metaphor, comparing neurodegeneration to communication system wires being 'cut at random'.[128] Brodman went on to explicitly trace out the causal pathway linking his experience of multiple sclerosis to his interest in formalising human reasoning: 'It was my observation of the subtle progressive changes in my higher functions that led me to

[126] Brodman to Howard Rusk, 26 Aug. 1954, KBP/7/6.
[127] Brodman to Norbert Wiener, 26 Apr. 1953, KBP/7/6.
[128] Brodman to Bigelow, 15 Apr. 1955, KBP/3/11.

believe the basic unit of these functions to be a simple, similar in many areas, and describable in mathematical terms'.

Brodman's letters to Wiener and Bigelow point to the pathway by which Brodman, a clinician trained in psychiatry and psychosomatic medicine, came to work on computing projects. Wiener and Bigelow were founders of the field of cybernetics.[129] In 1948, a year before Brodman and his colleagues released the CMI, Wiener published his influential book *Cybernetics.*[130] The book defined the field of cybernetics, the study of 'control and communication in the animal and the machine', and helped establish some of its guiding tools (mathematical models) and metaphors (mind as machine). Though peppered with complex equations and technical language, *Cybernetics* aroused considerable interest among both academic and lay audiences. The book's content (and title) struck most of its readers as utterly new, but the book in fact popularised, synthesised, and elaborated upon a line of thinking that had been in the making for a few years. In 1943, the neurophysiologist Warren McCulloch, together with the twenty-year-old logician Walter Pitts, published a landmark paper that described the first mathematical model of a neural network. The paper, demonstrating that neural nets could perform the logical operations of a Turing machine, invigorated ideas that mind and machine were fundamentally analogous.[131]

These early cybernetic ideas appealed to scholars with a wide range of disciplinary backgrounds—from the anthropologist Margaret Mead to the

---

[129] On Wiener, Bigelow, and cybernetics, see Peter Galison, 'The Ontology of the Enemy: Norbert Wiener and the Cybernetic Vision', *Critical Inquiry* 21 (1994): 228-266.
[130] Norbert Wiener, *Cybernetics: Or, Control and Communication in the Animal and the Machine* (New York: J. Wiley, 1948).
[131] On McCulloch and Pitts, see Tara H. Abraham, *Rebel Genius: Warren S. McCulloch's Transdisciplinary Life in Science* (Cambridge and London: MIT Press, 2016).

mathematician and engineer Claude Shannon. Between 1946 and 1953, this diverse network of researchers came together for a series of Cybernetics Conferences, organised by the Josiah Macy Jr. Foundation.[132] The conferences served as a venue for thinking through cybernetics in its many different instantiations: as a mathematical theory of communication and control, as a transdisciplinary study of the analogies between minds and computers, as an engineering enterprise concerning the construction of automata, and even as a more philosophical inquiry into what it means to be human.[133] Though heterogeneous, cybernetics animated widespread hopes that the behaviours of both inanimate and living systems could be modelled mathematically. If the Macy Conferences served as the launchpad for the field of cybernetics, the Dartmouth Summer Research Project on Artificial Intelligence, a workshop held in 1956, helped to introduce artificial intelligence as a field. The aim of the Dartmouth workshop, organized by the mathematician John McCarthy, was described in a proposal written for the conference in 1955: 'The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it'.[134] Through the 1960s, the 'symbolic artificial intelligence' approaches that came out of the Dartmouth workshop moved to the mainstream while overtly 'cybernetic' approaches migrated increasingly to the fringes.

---

[132] For a gripping account of the Macy Conferences, see Steve J. Heims, *The Cybernetics Group* (Cambridge: MIT Press, 1991).
[133] This variability has led the historian Ronald Kline to refer to the 'disunity' of cybernetics. Ronald Kline, *The Cybernetics Moment, Or Why We Call Our Age the Information Age* (Baltimore: Johns Hopkins University Press, 2015).
[134] John McCarthy, Marvin Minkey, Nathaniel Rochester, Claude Shannon, 'A Proposal For the Dartmouth Summer Research Project on Artificial Intelligence', *AI Magazine* 27 (2006): 12-14, on p. 12. This is a reprint; the original is dated 31 August 1955.

While cybernetics is often conceived of as a militaristic science, the sociologist Andrew Pickering has helped elevate the role of psychiatry in the field's emergence and development. As Pickering describes it, psychiatry was a 'surface of emergence' for the field of cybernetics—a disciplinary arena in which cyberneticists thought through ideas about feedback, control, and communication.[135] These intellectual exchanges between psychiatry and cybernetics arose out of a shared interest in the human nervous system. A central aim of cybernetics was, as Wiener noted in a 1948 essay, to study 'the common elements in the functioning of automatic machines and of the human nervous system'.[136] And the hallmark approach within cybernetics for probing these 'common elements' between minds and machines was the development of mathematical models of the nervous system. Brodman's early musing about how to formalise clinical cognition therefore carry a distinctively cybernetic flavour. Like Wiener, Bigelow, and other early cyberneticists, Brodman exhibited not only an intense fascination with the analogies between the human mind and automatic systems, but also an inclination to pursue these analogies through mathematical modelling. Beyond showing *that* Brodman was attracted to cybernetic approaches, these exchanges with Wiener and Bigelow also suggest *why* Brodman was so drawn: the letters showcase his embodied sensitivity to cybernetic principles of communication and control. What primed Brodman to the possibility of applying mathematical models and computers to the problem of medical diagnosis and decision making was his own experience of illness—an alertness to the sensations of his own body.

---

[135] Andrew Pickering, *The Cybernetic Brain: Sketches of Another Future* (Chicago: University Chicago Press, 2010).
[136] Quoted in Kline, *The Cybernetics Moment*, p. 12.

Citing the weight of other commitments, Bigelow backed out of their collaboration a few months after it began.[137] But around the time of Bigelow's departure, in the summer of 1955, Brodman received a serendipitous letter introducing him to someone who might be able to fill Bigelow's shoes: Adrianus van Woerkom.[138] A Dutchman, van Woerkom trained in astronomy at Leiden University, earning his doctorate in 1948. That same year, van Woerkom began serving as an Assistant Professor of astronomy at Yale University, where he also directed the nascent Computing Laboratory. Van Woerkom wore many hats as an academic: he was introduced to Brodman as 'Adrianus J.J. van Woerkom, astronomer, mathematician, and master of the IBM machine'. One of van Woerkom's interests was the application of mathematics to psychology—particularly areas of psychology where there existed large amounts of high-quality data. This was precisely the kind of area in which Brodman believed himself to be working. Human psychology was Brodman's chief intrigue, and he was working with large quantities of patient data procured from the CMI beginning in 1948. Brodman therefore followed up with van Woerkom presenting his research 'problem' as he had presented it to Oppenheimer some seven months earlier. Within a month, the two began a collaboration that would become the most important of Brodman's career.

In his collaboration with van Woerkom, Brodman continued his cybernetically inflected quest to develop a model of human cognition, one that might serve as the basis for his computerised diagnostic system. Contemplating the implications of such a model often left Brodman palpably excited. In 1960, as he was brainstorming a new model, he wrote to van Woerkom, 'I have so many thoughts about the new model that I

[137] Brodman to Bigelow, 23 Jul. 1955, KBP/3/11.
[138] Charles Robinson to Brodman, 1955, KBP/6/3.

feel like Handel when he said after he finished an oratorio that it was as if all heaven was opened before him. I feel as if all truth wer[e] opened before me'.[139] At this early stage of development, Brodman's ideas were both vague and grandiose, and they centered on one primary topic: pattern recognition. Diagnosing disease, Brodman ventured, was fundamentally an exercise in discerning patterns. As he expounded to van Woerkom, 'A physician certainly does not make computations when he makes a decision, nor does a person consciously use statistical techniques when he makes any decision. People recognize and manipulate patterns, gestalt. The model will do likewise'.[140] In the context of medical diagnosis, Brodman ventured that a pattern might consist of a 'configuration of symptoms'. It was the whole collection of symptoms that served as the meaningful unit of analysis in diagnostic reasoning. As Brodman described it, a physician, after 'recogniz[ing] some elements of a disease concept' would 'search for other elements [of the pattern] to convince himself that the disease concept exists in reality'.[141] According to Brodman, that diagnosing disease involved pattern recognition was simply a consequence of the fact that pattern recognition underlay all human (even animal) cognitive processes:

> The principles of the model are by no means limited in application to the making of a diagnostic decision. What I like most about the model is that we use a single simple principle, that the CNS recognizes reality by its pattern, to explain decision function, psychiatric phenomena, thinking, animal behavior, and many other living activities that heretofore have required complicated and unrelated formulations to explain each independently. It is not a large step to extend this principle to explain how the CNS recognizes abstract phenomena such as mathematical concepts, beauty, concept formation, creativity, and language, and how the CNS creates.[142]

[139] Brodman to van Woerkom, 3 Feb. 1960, KBP/6/3.
[140] Brodman to van Woerkom, 31 Jan. 1960, KBP/6/3.
[141] Brodman to van Woerkom, 3 Feb. 1960, KBP/6/3.
[142] Brodman to van Woerkom, 4 Feb. 1960, KBP/6/3.

Brodman wasn't just computerising diagnosis; by his estimation, he was modelling a unified theory of cognition.

These reflections on pattern recognition bear the mark of contemporaneous research in the fields of cybernetics and artificial intelligence. Pattern recognition emerged as a favourite topic within cybernetics owing in large part to Warren McCulloch and Walter Pitts' 1947 article 'How We Know Universals: The Perception of Visual and Auditory Forms'. Aiming to shed light on the neural basis of 'the *Gestalten* of Wertheimer and Köhler', the article outlined two neural mechanisms 'which exhibit recognition of forms'. That is, the article was an effort to understand (and describe mathematically) how humans could, for example, perceive a given object as the same (as 'invariant') even when viewing it from different sides and angles.[143] Related discussions of gestalt perception also figured prominently in *Cybernetics*.[144] By the late 1950s, the intellectual seeds of pattern recognition research that had been planted within the field of cybernetics began to germinate within the field of artificial intelligence.[145]

Pattern recognition offered one important cybernetic analogy between digital computers and human minds; psychiatric disorder offered another. Various cyberneticists drew parallels between psychiatric problems and problems in the operation of a computer. Humans develop psychoses and neuroses; machines crash, stall, and run amiss. Even psychiatric treatments had their digital correlates: Wiener believed, for instance, that electroshock therapy was akin to the wiping of a computer's

---

[143] Warren S. McCulloch and Walter Pitts, 'How We Know Universals: The Perception of Auditory and Visual Forms', *Bulletin of Mathematical Biophysics* 9 (1947): 127-147, pp. 127-128.
[144] Wiener, *Cybernetics*, p. 133.
[145] See, e.g., Oliver Selfridge, 'Pandemonium: A Paradigm of Learning', in National Physical Laboratory, *Mechanisation of Thought Processes* (London: Her Majesty's Stationary Office, 1959), pp. 511-529.

memory.[146] Here again, Brodman's views are remarkably consonant with those of early cyberneticists. Brodman mused that an ideal model would be one that failed in the same way that human minds tend to fail. He wanted the model to be a little bit 'neurotic': 'In general', Brodman wrote in some notes in 1955, 'the math model should have the possibility of doing what humans do, even of developing the same defects in judgement and behavior as a human with a neurosis or psychosis'.[147] It would be a mistake to claim that Brodman was a full-fledged cyberneticist; he wasn't. Brodman never cited any key cybernetic works in his writings, nor did he identify his views as being cybernetically inspired.[148] But in light of his unsolicited letter to Norbert Wiener, his engagements with Julian Bigelow, and the parallels between cybernetic thought and his own, I suggest that Brodman's attraction to certain threads of cybernetic thinking helps explain how a project of medical history taking began to evolve into one of computerising diagnosis.

THE CLINICIAN AND THE CLERK

Brodman was not the only actor theorising about the possible mathematical basis of the medical mind. Nor was all of this theorising pursued within a cybernetic framework. How much does statistical inference resemble clinical inference? In the

---

[146] Kline, *The Cybernetics Moment*, p. 12.
[147] Untitled notes, 20 Apr. 1955, KBP/3/11. No author is listed for these notes. Brodman's authorship can be inferred with a considerable level of confidence from the notes' content and style, which accord with that of his other writings, correspondence, and notes. See, e.g., Brodman to van Woerkom, 12 Oct. 1957, KBP/6/3.
[148] Brodman did, however, read materials that were rich in cybernetic thought. In 1960, he wrote to van Woerkom, 'I am sending you the books on *Mechanisation of Thought Processes*. You will find a great deal of material in it of interest to us'. Brodman to van Woerkom, 5 May 1960, KBP/6/3. These volumes were the published versions of a symposium held at the National Physical Laboratory in November of 1958. The symposium brought together leading figures in the fields of cybernetics, artificial intelligence, pattern recognition, and other related areas.

mid-twentieth century, a vigorous debate materialised around this question, particularly within the realm of clinical psychology. The question of whether clinical and statistical modes of reasoning were similar soon shaded into the more charged question of which was superior. One researcher characterised the controversy in a 1963 article: 'The much-debated issue at hand seems to be whether judgment based on broad experience, sound inference, and a reflective weighing of alternative points of view can achieve or exceed the accuracy of prediction based on quantification of relevant aspect of present and/or past behavior'.[149] The researcher went on to further describe the opposing camps: 'Those who incline toward the quantitative point of view argue that the human intellect is incapable of storing, combining, and weighing information with the efficiency of a statistical formula tailored for use in an electronic computer'.[150] By contrast, '[t]hose who have greater faith in judgements of the human mind are usually willing to concede the superior efficiency of a numerical statement, but would point out that only limited aspects of complex problems in which human behavior is the focus of study can be easily quantified. If formal quantitative reasoning can only partially solve a problem, then qualitative assessment, based on supplementary nonquantified information, intuition, and experience, must be introduced to fill the void'.[151]

Among the more important participants in this debate was the young clinical psychologist Theodore Sarbin. Sarbin's academic journey to the topic of clinical and statistical reasoning was a circuitous one. Raised in poverty, Sarbin worked a number of odd jobs during his youth, completing secondary school requirements at night. After

[149] Jesse G. Harris, Jr., 'Judgmental versus Mathematical Prediction: An Investigation by Analogy of the Clinical versus Statistical Controversy', *Behavioral Science* 8, no. 4 (1963): 324-335, p. 324.
[150] Ibid., p. 324
[151] Ibid., p. 325.

the onset of the Great Depression in 1931, he 'road the rails' as a vagabond.[152] On a twenty-dollar loan Sarbin attended Ohio State University, and by the time he completed his graduate work there in the early 1940s Sarbin had established himself as something of an intellectual maverick. Some of Sarbin's earliest publications, based on his doctoral work, concerned the connection between clinical and statistical ('actuarial') prediction. In these early writings, Sarbin took the unconventional position that the predictions of clinicians were, when stripped to their core, actuarial.[153] Like actuarial methods of prediction, which rely on known correlations between symptoms and outcomes, clinical prediction, Sarbin claimed, boiled down to the probabilistic linking of the present case to impressions of previous, comparable cases. The actuarial character of clinical prediction is apparent when clinicians talk about their patients' 'chances', or when describing a certain outcome as 'quite likely'. The difference, then, between these two methods is a matter of degree: clinical prediction mimics actuarial prediction, only the clinical side of things is also subject to the 'human factor'—the errors of the human mind. According to Sarbin, any divergence between the two methods reflects the crud of human error gumming up the machinery of clinical prediction. The clinician cannot in principle outperform the actuary.

In 1944, Sarbin's work caught the eye of the young psychologist Paul Meehl. Meehl encountered Sarbin's research at a time when he was intellectually wrestling with competing ideas about clinical practice. Trained in the psychoanalytic tradition, Meehl was sympathetic to the psychoanalytic enterprise and its clinical aims. But in the

---

[152] Leslie Dunn, 'The Wander Years', *Monterey County Herald*, 17 Dec. 2004.
[153] Theodore Sarbin, 'Clinical Psychology—Art or Science?', *Psychometrika* 6 (1941): 391-399; Theodore Sarbin, 'The Relative Accuracy of Clinical and Statistical Predictions of Academic Achievement', *Psychological Bulletin* 38 (1941): 714; Theodore Sarbin, 'A Contribution to the Study of Actuarial and Individual Methods of Prediction', *American Journal of Sociology* 48 (1943): 593-602; Theodore Sarbin, 'The Logic of Prediction in Psychology', *Psychological Review* 51 (1944): 210-228.

1940s, Meehl also found himself fascinated by the views of more empirically inclined psychologists and psychiatrists who expressed scepticism toward Freudian ideas, particularly those based on 'clinical experience' rather than statistical data. The result of these competing intellectual commitments was an 'intense cognitive dissonance and a strong, persistent need to resolve it'.[154] Resolution, for Meehl, took the form of a decade-long examination of the relationship between clinical and statistical methods of prediction. In 1954, he collected and published his thoughts on the subject in the landmark 1954 book *Clinical versus Statistical Prediction*.

Meehl wrote the book in part as a critique of Sarbin's perspective. In contrast to Sarbin, Meehl believed in clinician exceptionalism—he believed that clinicians were able to do things with clinical information 'beyond that which can be done by the mechanical application of an actuarial table or a regression equation'.[155] Meehl, himself a clinician, paid his dues to what he called the 'special powers of the clinician'. To illustrate his point, Meehl marshalled a clinical case study reported by the psychoanalyst Theodor Reik in his famous book *Listening with the Third Ear*.[156] As related by Reik, a patient came to see him complaining of a toothache. Reik listened as patient discussed the toothache and her recent trips to the dentist, before falling into period of silence. Then, breaking the silence, she pointed to Reik's bookcase and said, 'There's a book standing on its head'.[157] This combination of experiences, symbols,

---

[154] Paul E. Meehl, 'Paul E. Meehl', in G. Lindzey (ed.), *A History of Psychology in Autobiography* (Stanford: Stanford University Press, 1989), 337-389, p. 359.
[155] Paul E. Meehl, *Clinical versus Statistical Prediction: A Theoretical Analysis and Review of the Evidence* (Minneapolis: University of Minnesota Press, 1954), p. 24. Meehl's critique of Sarbin's position also included various technical points which fall outside the scope of this dissertation. For an analysis of these, and for other details on Meehl's contribution to the topic, see William Grove, 'Clinical versus Statistical Prediction: The Contribution of Paul E. Meehl', *Journal of Clinical Psychology* 61, no. 10 (2005): 1233-1243.
[156] Theodor Reik, *Listening with the Third Ear* (New York: Farrar & Straus, 1948).
[157] Quoted in Meehl, *Clinical versus Statistical Prediction*, p. 51.

and clinical cadence prompted Reik to ask his patient, 'without the slightest hesitation', 'But why did you not tell me that you had had an abortion?' Examples of clinical reasoning like these left Meehl appreciating the limits to actuarial approaches. How could crude actuarial law hope to capture and formalise the kinds of interpretations, generalisations, and abstractions that clinicians like Reik sometimes masterfully employed? Say what you will about the clinician's first two ears. But, to Meehl, the proverbial 'third year' seemed to fall clearly outside of the actuary's remit.

Despite the lengths Meehl went to distinguish his views from those of Sarbin, his book was ultimately discussed and remembered for the occasions on which their views converged. To claim, as Meehl did, that clinical prediction differs from actuarial prediction was not to claim the superiority of the former over the latter. In fact, Meehl agreed with Sarbin that statistical, and otherwise mechanical, methods produced better and more efficient decisions about medical prognosis and treatment. This, for Meehl, was an empirical question. After reviewing twenty empirical studies on statistical and clinical prediction, Meehl concluded bluntly, 'it is clear that the dogmatic, complacent assertion sometimes heard from clinicians that "naturally," clinical prediction, being based on "real understanding," is superior [to statistical prediction], is simply not justified by the facts collected to date'.[158] Meehl saw in statistical prediction tools a precision, reliability, and objectivity that was often lacking in clinicians. To drive home his broader point about the importance of statistics and actuarial methods in clinical medicine, Meehl entertained the possibility that he was wrong—that clinical predictions were in fact superior to actuarial ones. Meehl pointed out, however, that even if the accuracy of the clinician surpassed that of the statistician, this difference

[158] Meehl, *Clinical versus Statistical Prediction,* p. 119.

must be shown statistically. Put simply, statistics in medicine is unavoidable: '*always* the actuary will have the final word'.[159]

Many of the nuances of Meehl's positions were lost in the controversy the followed. 'The controversy surrounding this issue', wrote on clinical psychologist, 'has been intense'.[160] The intensity of the debate might have been predicted: the notion that rote statistical prediction edged out artful clinical prediction offended the sensibilities of clinicians who believed clinical decision making to be more of an art than a science. Benjamin Kleinmuntz sketched out this view in a 1968 article:

> Clinical intuition, it is often said, is an art and as such it cannot be subjected to the rigors of scientific study. Clinicians tenaciously cling to the opinion that patients communicate with, or single to, them in numerous nonverbal modes, and they tend to believe that they must see, and sometimes even "smell out," their patients' diseases in order to arrive at an accurate diagnosis. Some clinicians, especially in the psychiatric areas, report that they "resonate" to patients with their unconscious and claim to gain some special understanding of persons' motivations and dynamics this way. Other clinicians believe that they possess special power of intuition which permit them to find hidden meanings beneath even the most neutral gestures.[161]

Twenty years later, a bioethicist also put words to the sentiments driving much of the disquietude towards Sarbin and Meehl's work: 'One need not be a mystical intuitionist to argue that, so far, at least, it seems that formalisms alone will not lead to what could be described as intelligent behavior. A clinician gains understanding and clues from a patient's choice of words, his facial expressions, his gestures, his movements, and his subliminal messages. The clinician', he continued, 'can empathize with the patient's

---

[159] Ibid., p. 138.
[160] Richard D. Mann, 'A Critique of P. E. Meehl's *Clinical versus Statistical Prediction*', *Behavioral Science* 1, no. 3 (1956): 224-230, p. 224.
[161] Benjamin Kleinmuntz, 'The Processing of Clinical Information by Man and Machine', in Benjamin Kleinmuntz (ed.) *Formal Representation of Human Judgment* (New York: John Wiley, 1968).

emotional state and come to see how the patient is interpreting his illness'.[162] (Meehl, it should be noted, likely would not have disagreed with this view.) The art of French cooking may have been amenable to a masterful cookbook; Julia Child, however, may not have done so well with the art of medicine.

Other criticisms of Meehl's stance were more measured—some dealing with technical matters of mathematics, others with technical matters of philosophy. A psychologist at the University of Michigan, Paul Mann, took issue with the manner in which Meehl compared statistical and clinical modes of reasoning. Mann charged that the tasks on which the two modes of reasoning were compared did not adequately capture the complexities of the clinical environment. When clinicians make decisions, their task goes beyond the mere estimation of immediate outcomes. Clinical decisions, Mann argued, are rationally shaped by many other social, cultural, and personal values that were not adequately reflected in either Meehl or Sarbin's accounts. Mann explained, 'If we are to understanding why someone bets on this treatment rather than that, it is important to keep in mind that not only the probability of cure is relevant but that estimated losses and gains resulting from a decision must also be considered'.[163] Mann brought concreteness to his point with an example: 'If, for example, studies showed that this year 3 per cent of all seven-year-old boys were going to die from polio and that if a certain vaccine were used, none would die from polio but 3 per cent would die as a direct result of the vaccination, it seems a reasonable guess that the parents (and probably the doctors) would oppose substituting man-made death for natural death. The consequences of the act of leaving the boy to a more mysterious agent of

[162] K. Danner Clouser, 'Approaching the Logic of Diagnosis', in Kenneth F. Schaffner (ed.) *Logic of Discovery and Diagnosis in Medicine* (Berkeley, CA: University of California Press, 1985), 35-55, p. 52.
[163] Mann, 'A Critique of P. E. Meehl's *Clinical versus Statistical Prediction*', p. 229.

death are obviously greatly different from the consequences of doctor-produced illness'. To compare the clinician and the actuary solely on the basis of some immediate and narrow outcome reproduced a feeble understanding of real-world medical decision making. Of course, Mann suggested, clinical judgements would deviate from actuarial judgements in prediction; clinical judgments, after all, weren't exclusively acts of prediction—they were doing other kinds of social, moral, legal, and cultural work.

The exchange between Meehl and Sarbin—and the larger discussions around them—represents only a small cross-section of the medical community's general interest in the mathematical nature of clinical reasoning. Through the mid-twentieth century, numerous other physicians and researchers also discussed the statistical and logical dimensions of clinical prediction, reasoning, and judgement. The mathematician John Davis Williams, for instance, suggested that the newly developed techniques of game theory be brought to bear on medical decision making.[164] Along similar lines, the psychologist W. K. Estes proposed mathematical models for clinical diagnosis in his widely discussed presidential address to the Midwestern Psychological Association. Estes mused in the 1957 address, 'we might indeed live in a world where all natural phenomena yield alike to the analytical weapons of mathematics and experiment.'[165] By 1959, Robert Ledley and Lee Lusted opened the doors to discussions about the Bayesian foundations of clinical reasoning, as the following section of this dissertation will describe.[166] Collectively these writings indicate that, in the years leading up the

---

[164] John D. Williams, *The Compleat Strategyst: Being a Primer on the Theory of Games of Strategy* (New York: McGraw-Hill, 1966 [1954]), pp. 122-25.
[165] W.K. Estes, 'Of Models and Men', *American Psychologist* 12, no. 10 (1957): 609-617, p. 617.
[166] Robert S. Ledley and Lee B. Lusted, 'Reasoning Foundations of Medical Diagnosis', *Science* 130, no. 3366 (1959): 9-21.

first efforts to computerise medical diagnosis, researchers increasingly discussed the possibility that diagnostic reasoning was (or could be) supported by a mathematical scaffold.

Running through many of these discussions were allusions to 'Hollerith machines', electromechanical tabulating machines that analysed information recorded on punched cards. Often images of early computational devices were not invoked out of any genuine interest in computerising clinical reasoning; they were invoked, rather, for rhetorical effect. The Hollerith machine came to symbolise Sarbin's view in the extreme: that clinical reasoning is no different from statistical reasoning. For example, Meehl suggested that Sarbin's viewpoint could be reduced to the simple statement that 'the clinician is a second-rate substitute for a Hollerith machine'.[167] But within a few years, references to the Hollerith machine, and to other early computational tools, in discussions of clinical reasoning no longer served a predominately rhetorical function. By the mid-1950, people like Brodman hazarded that the task of medical diagnosis might actually be well-suited to the computer—and they were taking concrete steps develop computerised methods of diagnosis.

Having long been engaged in 'armchair considerations' about when, and under what circumstances, human and statistical judgements diverged, Meehl was, by late 1959, ready to enter the actual playing field.[168] One of the playing fields on which Meehl tested clinical versus statistical prediction was the Minnesota Multiphasic Personality Inventory (MMPI)—a standardised psychological test that assessed

[167] Meehl, *Clinical versus Statistical Prediction,* p. 76.
[168] Paul E. Meehl, 'A Comparison of Clinicians with Five Statistical Methods of Identifying Psychotic MMPI Profiles', *Journal of Counseling Psychology* 6, no. 2 (1959): 102-109, p. 102.

personality and psychopathology.[169] The MMPI's name hinted at the circumstances around its creation. The MMPI came into being during the late 1930s and early 1940s at the University of Minnesota. Its principal creators—the clinical psychologist Starke R. Hathaway and the neuropsychiatrist J. C. McKinley—described the impetus behind its development. The development and clinical evaluation of 'an extensive questionnaire method of measuring symptomatic patterns in the personality', McKinley described in a 1941 letter, would serve two key objectives: first, it would 'provide a more objective measure of the relative strength of such symptomatic factors as hypochondriasis, depression and hysteria', and second, it would serve as a 'powerful instrument against which the opinions of young men in training for psychiatry can be checked'.[170] In a field dominated by long unstructured interviews, subjective professional appraisals, and unstable value judgments, a tool like the MMPI, its promoters hoped, could bring a more structured and objective approach to the assessment and classification of human personality.

As originally published in 1943, the MMPI featured 504 yes-no questions which had been selected from an original pool of over one thousand inventory items brainstormed by Hathaway and McKinley. The researchers then compared questionnaire item responses among healthy subjects ('Minnesota normals') to item

---

[169] On the history and development of the MMPI, see, e.g., Roderick D. Buchanan, 'The Development of the Minnesota Multiphasic Personality Inventory', *Journal of the History of the Behavioral Sciences* 30 (1994): 148-161; Roderick D. Buchanan, 'Ink Blots or Profile Plots: The Rorschach versus the MMPI as the Right Tool for a Science-Based Profession', *Science, Technology, and Human Values* 22, no. 2 (1997): 168-206; Rebecca Schilling and Stephen T. Casper, 'Of Psychometric Means: Starke R. Hathaway and the Popularization of the Minnesota Multiphasic Personality Inventory', *Science in Context* 28, no. 1 (2015): 77-98; Deborah J. Coon, 'Standardizing the Subject: Experimental Psychologists, Introspection, and the Quest for a Technoscientific Ideal', *Technology and Culture* 34, no. 4 (1993): 757-783.

[170] J.C. McKinley to Theodore C. Blegen, 19 May 1941, Starke Rosencrans Hathaway Papers, University Archives, University of Minnesota, Minneapolis, Minnesota, box 1, folder 14 (hereafter cited in the form SHP/1/14).

responses among subjects with known psychiatric or personality disorders. From these data, Hathaway and McKinsey constructed ten clinical dimensions (scales) along which to plot a subject's personality structure (e.g, 'hypochondiasis', 'schizophrenia', 'paranoia', 'depression', etc.). Meehl played a hand in the MMPI's creation as well. In a 1969 letter to Hathaway, Meehl half-facetiously characterised the contributions of the three men as follows: 'It is perhaps a worthwhile half-truth to say of the MMPI, "McKinley needed it, Hathaway built it, Meehl expounded it"'.[171] Specifically, Meehl developed, as part of his doctoral dissertation under Hathaway, the K scale for detecting the honesty of subjects' responses. Meehl's intellectual maturation at Minnesota, working on the MMPI's creation, fed directly into his interest in actuarial approaches to clinical decision making. As Meehl later wrote to Hathaway, in a letter thanking him on the occasion of his retirement, 'And for my own nefarious purposes, how far would I have got applying my actuarial, cookbook ideas if I had been saddled by a mentor (or school) with an addiction to inkblocks or the TAT?'[172]

It made sense that the MMPI came out of Minnesota. While other psychiatry and clinical psychology departments at the time organised themselves around large theoretical approaches (Freudian psychoanalysis, Skinnerian behaviourism), Minnesota hewed instead to a streak of eclectic, moderate, and atheoretical empiricism. Meehl characterized Minnesota as the guardian of the 'the tough-minded, theory-skeptical, "dustbowl empiricist," *pragmatic* tradition'.[173] The MMPI's creation fed back into this reputation.

Because of the MMPI's constrained and standardised nature, a number of different 'techniques', 'codes', and 'rules'—i.e., statistical models—were created to

[171] Paul Meehl to Starke Hathaway, 19 Mar. 1969, SHP/1/20.
[172] Ibid.
[173] Ibid.

guide MMPI profile interpretation.[174] Meehl put these models to the test. In 1959, he published a study on the performance of five such statistical models, comparing them to the performance of human clinicians.[175] The findings were stark. The three 'pure actuarial' models, Meehl discovered, outperformed the individual clinicians in identifying psychotic patients. A string of experimental findings like this informed Meehl's broader perspective on the use of statistical methods in the clinical setting.

On 29 April 1955, Meehl delivered, in his iconic crisp and punchy style, the presidential address at the Midwestern Psychological Association annual meeting in Chicago, IL. The address's title was a playful and provocative nod to Meehl's critics: 'Wanted—a good cookbook'.[176] Its content was a strong argument for the application of 'cookbook' (i.e., statistical, actuarial) methodologies to the interpretation of MMPI profiles. Meehl clearly believed that actuarial methods of MMPI interpretation would take the field far. But he cautioned the overzealous with the same sharpness and verve with which he criticised the underenthusiastic: 'Statisticians (and rat men) with castrative intent toward clinicians should beware of any temptation to overextend these findings that "clinicians don't actually add anything"'.[177] Evidently Meehl's words of caution were not well heeded; his address tapped into a worrisome emotionalism on both sides. As Meehl later recalled,

> My Midwestern Presidential Address led to an episode which puzzled and troubled me at the time as reflecting a serious problem in the profession.

---

[174] See, e.g., D.T. Lykken, 'A Method of Actuarial Pattern Analysis', *Psychological Bulletin* 53 (1956): 102-107; E.C. Taulbee and B.D. Sisson, 'Configurational Analysis of MMPI Profiles of Psychiatric Groups', *Journal of Consulting Psychology* 21 (1957): 413-417; P.E. Meehl and W.G. Dahlstrom, 'Objective Configural Rules for Discriminating Psychotic from Neurotic MMPI Profiles', *Journal of Consulting Psychology* 24 (1960): 375-387.
[175] Meehl, 'A Comparison of Clinicians with Five Statistical Methods of Identifying Psychotic MMPI Profiles'.
[176] Paul E. Meehl, 'Wanted—A Good Cookbook', *American Psychologist* 11, no. 6 (1956): 263-272.
[177] Meehl, 'Wanted—A Good Cookbook', p. 263.

> Though aware of the tension between clinical practitioners and academic experimental psychologists, I was surprised by its emotional intensity and was not skillful at defusing it. In presenting empirical data relevant to the idea of formalizing profile interpretation rather than "clinical eyeballing," I had told a couple of funny stories, employing some snide expressions about clinicians who reject objective data. Following the talk, which was well received both by scholarly clinicians and nonclinicians, I was invited for drinks in the hotel room of a distinguished experimental psychologist. There were a half-dozen of his experimental brethren along with two academic clinicians. The general flavor of the discussion was "Meehl, you sure gave those clinicians a good beating," an overinterpretation of my message which I found troublesome but let pass. The sentiment was that it was fine to see a clinical psychologist who also ran rats and knew how to take a partial derivative getting elected to a prestigious office and thereby provided with a big audience. The experimentalists had not seen my recent book on clinical and statistical prediction, but via the *anti*-actuarial arguments in that book, one of the clinicians was able to bring up the subject of the clinician's "third ear" and those kinds of inferences about the psychodynamics or historical past that it would be hard to imagine objectifying.[178]

The 'cookbook' approaches outlined theoretically by Meehl adumbrated the practical application of computers to MMPI interpretation. Less than a decade after the publication of Meehl's lighting-rod book, the first computerised version of the MMPI was developed and implemented at the Mayo Clinic in Rochester, Minnesota. The initiative was less a preconceived effort than the product of happy circumstance. It was, a user's guide noted, 'a fortunate combination of circumstances' that 'brought together a highly motivated professional staff, a group skilled in electronic data processing, and a large population of receptive medical patients'.[179] 'This combination', the report continued, 'gave us the opportunity to implement the *practical application* of computerized personality assessment procedures in a medical setting'. Part of the luck

[178] Paul E. Meehl, in G. Lindzey (ed.), *A History of Psychology in Autobiography* vol. 8 (Stanford: Stanford University Press, 1989), 337-389, pp. 353-354.
[179] Wendell M. Swenson, David Osborne, Robert C. Colligan, 'A User's Guide to Mayo Clinical Computerized Scoring and Interpretive System for the Minnesota Multiphasic Personality Inventory (MMPI), Second Edition', p. 1, Pamphlet Collection, box 04, folder 'User's Guide – Mayo Clinic MMPI', Mayo Clinic Historical Unit, The Mayo Clinic, Rochester, Minnesota (hereafter cited in the form Mayo/Pamphlet Collection/04/'User's Guide – Mayo Clinic MMPI').

also derived from a happy coincidence of location. In 1958, IBM moved its Midwest branch to Rochester. Members of Mayo's Section of Psychology observed the move with intrigue. As Mayo's chief of psychiatry recalled, 'When IBM moved its Midwest plant here in 1958, some of us went out to find out what these computers were all about. And as a consequence of then being surfeited with requests for off-the-floor consultations, we decided among us that we would devise a plan whereby each consultant at the Mayo Clinic could be acquainted with a device that could tell him the kind of patient that he was treating so that the selection process for subsequent referral to psychiatry would be enhanced thereby. And so then the automated version of the MMPI came into vogue …'.[180] By 1961, the programme was operational.

The Mayo Clinic's internal newspaper, *MAYOVOX*, announced the computerised MMPI initiative to employees in a February 1962 article. 'A pilot study of the use of a "personality" test as part of the general physical examination and the processing of such tests by IBM machines has been carried out at the Clinic during recent months'.[181] The article described the method's operation. MMPI tests completed by patients were delivered to Mayo's Section of Systems and Procedures, where they were coded on IBM cards. These cards were then outsourced to an IBM facility for scoring with a 650 computer. 'The marked cards are fed into the computer, where, by means of a combination of magnetic tapes and magnetic core storage, operations are carried out to score the individual record on thirteen clinical and validity scales. The computer is programmed to select from a statement library of 72 statements those which best describe the patient's personality based on his MMPI responses'.[182] Such a

[180] Howard P. Rome, staff memories, p. 12, Mayo/MHU-0670/022/F393.
[181] 'Written "Personality Tests" Promising As Clinical Tool', *MAYOVOX* 13, no. 7 (3 Feb. 1962): 1, 3, p. 1, Mayo/Mayovox vol. 13.
[182] 'Automated MMPI Proves Useful Tool For Clinical Diagnosis and Research', p. 2.

personality statement might read: 'Rigid and meticulous. Dissatisfied with social relationships … moderately depressed, worrying and pessimistic … sensitive .. alive to opinions of others'.[183] Though this was an involved procedure, the speed of machine analysis left its human counterparts looking plodding by comparison. The *MAYOVOX* article boasted, 'The machine performs this operation in about 35 seconds. Approximately 45 minutes would be required to perform the same operation when a trained individual scored the cards and a psychologist interpreted the results and dictate a report'.[184] The speed and convenience of the process only improved after Mayo had opened its own computer centre, allowing for in-house processing of the MMPI tests.

While Mayo represented the first centre to computerize the MMPI, it was far from the only. In the early 1960s, Raymond Fowler, the chair of psychology at the University of Alabama, initiated a research enterprise to develop a computer programme for MMPI interpretation. Fowler described this effort in article provocatively subtitled, 'The Automated Psychologist': 'The goal was to produce a valid and clinically useful report in an easy to read narrative style'.[185] Fowler's research caught the attention of folks at Roche Laboratories, a subsidiary of the healthcare giant Hoffman-La Roche. Intrigued by the potential commercialisation of a computerised MMPI service, Roche officials reached out to Fowler and his colleagues to arrange a pilot test. To coordinate and oversee the pilot study, Roche created the Roche Psychiatric Service Institute (RPSI).

The pilot test proceeded in four phases. In the first phase, Roche affiliates interviewed psychiatrists, clinical psychologists, and representatives of relevant

[183] 'Written "Personality Tests" Promising As Clinical Tool', p. 3.
[184] Ibid., p. 1.
[185] Raymond D. Fowler, 'Computer Interpretation of Personality Tests: The Automated Psychologist', *Comprehensive Psychology* 8, no. 6 (1967): 455-467, p. 457.

professional associations to gauge the field's general reactions to a computerised scoring and interpretation service. Where phase one tested the method in theory, phase two tested the method in practice. To carry out computerised MMPI interpretation, RPSI established a robust network of hardware (computers), software (programs), and wetware (humans), one that linked the individual offices of clinicians, an RPSI processing centre in New Jersey, and a computer centre in Alabama. Completed MMPI answer sheets from clinicians' offices were sent to the RPSI processing centre in New Jersey, where they were scanned by a 1232 optical scanner. After being read into a 1004 terminal, the data was sent, via Dataphone, to the 1004 terminal at the computer centre located in Alabama. There, the data was moved to a Univac 1107 computer and, finally, interpreted according to programmed decision rules. A final report of the computer's analysis was rendered and mailed to the referring clinician. This test arrangement, the researchers found, 'operated smoothly and the response was favorable'.[186] Having established the viability of this computerised network, phase three aimed to assess the method's scalability. Roche interpreted the results of the test as a green light: 'On the basis of the pilot tests', Fowler reported, 'a decision was made by Roche Laboratories to make the service available nationally at a nominal fee as a professional service to psychiatrists, clinical psychologists and appropriate agencies and institutions'.[187] In 1965, Roche made Fowler's computerised system commercially available on a national basis.

These efforts to computerise the MMPI stoked familiar flames. Many clinical psychologists prided themselves on the 'art' of MMPI interpretation. They believed strongly in their ability to 'eye-ball', on the basis of skill, experience, and intuition,

[186] Fowler, 'Computer Interpretation of Personality Tests', p. 465.
[187] Ibid., p. 466.

various MMPI personality configurations. For some, the introduction of computers to MMPI interpretation challenged this approach—and, by extension, their sense of self, professional value, and identity. Hathaway himself acknowledged that the MMPI's results had to be refracted through the lens of 'clinical intuition' for the tool to be clinically useful. In his address at a 1969 symposium on the MMPI, he explained, 'The MMPI … needs much "clinical interpretation" despite statistically significant validities for many diagnostic and predictive applications'.[188] While still defending the MMPI, Hathaway warned against the 'overzealous' application of 'seductive features of mathematical order' to certain areas of knowledge and humanity.[189]

At Mayo, the reception of the automated MMPI among physicians and patients was mostly positive—at least as reported by *MAYOVOX*: '70 per cent of 158 Clinic consultants who have had extensive experience with the automated MMPI agree that the test is a valuable adjunct to a general medical examination, that it frequently reveals unsuspected kinds or degrees of emotional problems, and that it offers various other advantages with no serious disadvantages'.[190] Yet the situation was not uniformly rosy. Suggesting the possibility of dissenting viewpoints, another internal news report noted, 'experience of men in our section forms a spectrum'.[191] Predictably, then, developers of computerised approaches to MMPI interpretation emphasised the continued importance of human clinical judgement, even after computers entered the scene: 'It has never been our intent to replace the interpretative expertise of the experienced medical

---

[188] Stark R. Hathaway, 'Where Have We Gone Wrong? The Mystery of the Missing Progress', Presented at the Fourth Annual Symposium on the Recent Development in the use of the MMPI, University of Minnesota, 10 Apr. 1969, pp. 2-3, SHP/1/12.
[189] Ibid., p. 23.
[190] 'Automated MMPI Proves Useful Too For Clinical Diagnosis and Research', *MAYOVOX* 16, no. 3 (4 Dec. 1964), p. 2, Mayo/Mayovox vol. 16.
[191] 'Written "Personality Tests" Promising As Clinical Tool', *MAYOVOX* 13, no. 7 (3 Feb. 1962): 1, 3, p. 3, Mayo/Mayovox vol. 13.

psychologist', Mayo officials assured in the user's guide.[192] Their quotes in *MAYOVOX* struck a similar tone: 'We have all regarded the information the test provides as a bonus … *in addition to, not instead of,* the usual insight which we get during a patient interview, either intuitively or by systematic history taking'.[193] Debates and discussions around the MMPI—and its automation—thus tracked closely to those that swarmed around the CMI.

DIGITISING DATA

At the same time that Brodman was engaged in airy philosophical meditations concerning various models, he was also consumed by the labour and monotony of more earthy data practices. The CMI data collected from New York Hospital patients between 1948 and 1949 provided Brodman, van Woerkom, and their collaborators with three basic kinds of information about some five thousand patients: their diagnosis as determined by the New York Hospital clinical staff, their responses to the CMI, and their demographic characteristics. Originally, this 'raw' data took the form of thousands of paper questionnaires. The data, in this form, exhibited a lot of what the historian Paul Edwards has called 'data friction'—the questionnaires were bulky, they took up physical space, they demanded great labour to sort through, and they could be (and were) physically lost and misplaced.[194] At some point, by the mid-1950s, this information had been abstracted and coded onto punched cards, and later magnetic tape, for computer processing. Although difficult to know with certainty, these data

[192] Swenson, Osborne, and Colligan, 'A User's Guide', p. 1.
[193] 'Written "Personality Tests" Promising As Clinical Tool', p. 3.
[194] Paul Edwards, *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming* (Cambridge: MIT Press, 2010), pp. 83-110. Brodman noted to van Woerkom in 1956, 'As I told you yesterday, the listing you gave me on Friday indicates that the missing 10% of CMIs in the 1948 sample were not lost in a random manner as I had hoped'. Brodman to van Woerkom, 13 Mar. 1956, KBP/6/3.

may have also existed in an intermediate, abstracted paper 'database' as well. Brodman hinted at this possibility in a 1970 letter, when he noted that the 'original volume' of 1948-1949 data had been lost: 'All data derived from the Cornell Medical Index of course remains in the public domain, but by this time all except published material have been lost. The original volume of data was enormous with over 23 thousand entries for one table alone'.[195]

By the mid-1950s, Brodman's team had articulated an overarching strategy for building a diagnostic model from this information. Their model would permit making diagnoses from CMI data 'on the basis of the frequency of each complaint in different diseases'.[196] In other words, they hoped to exploit the fact that patients with a given disease responded to questionnaire items in a generally distinctive way. With this general framework in place, Brodman and his collaborators turned to questions of mathematics. Although building a mathematical model was necessarily technical, Brodman believed that his clinical intuition would effectively guide their early work. He wrote to van Woerkom, 'My clinical experience suggests that the complexities of the problem will make many pitfalls and booby traps, and we must be alert and careful in planning if they are to be avoided. Fortunately, clinical experience can not only help avoid these traps but can also yield the intuitive approach to successful planning that is so important in studying a new area of knowledge'.[197]

Ultimately, however, developing the mathematical model had a lot less to do with theory and intuition than with laborious, trial-and-error, and sometimes manual, paper-based work. The earliest work toward the model, conducted with the help of the mathematicians Albert Erdmann and Jerome Deustchberger (and before van Woerkom

[195] Brodman to Gerald Fleischli, 2 Oct. 1970, KBP/5/1.
[196] Brodman to van Woerkom, 1 Dec. 1955, KBP/6/3.
[197] Brodman to van Woerkom, 21 Dec. 1955, KBP/6/3.

joined the team), involved an item analysis of CMI responses with respect to 25 common diseases. At least some of this early work was done manually. When Brodman wrote to van Woerkom about this stage of the project, he noted, 'My notes go back many years when I was getting Jerry Deutschberger to make a hand analysis, which I showed you and commented on the first day we met'.[198] Brodman referred to this early work as a 'pilot study'—a way of testing the viability of a mathematical diagnostic method. Elaborating on this pilot study elsewhere, Brodman reiterated the manual dimension of this early work and the stamina that it likely involved: 'Because it is operated by hand, this statistical method was slow and cumbersome'.[199] The tedium paid off. Their manual analyses proved (to them at least) that a statistical—and thus a computerised—method of making differential diagnoses using the CMI data was possible. Before Brodman and van Woerkom went about computerising diagnosis, then, there was first an effort to confirm the project's viability at least partially by hand.

Beyond highlighting the importance of manual forms of calculation, the so-called pilot study points to a second way in which the development of the computerised MDS evolved from, and rested on, paper data practices. Brodman and his colleagues worked out their computerised method using data that had been collected much earlier and for a different purpose. Brodman and his colleagues amassed large amounts of CMI data beginning in 1948 not with the hopes of eventually computerising diagnosis. At that time, their aim was more immediate: to better understand the CMI and its clinical value. Years before Brodman and van Woerkom started to probe this data, it formed the basis of a series of articles on the CMI. That the same exact data guided Brodman's project of computerising diagnosis underscores how the MDS was enabled

[198] Brodman to van Woerkom, 16 Dec. 1958, KBP/6/3.
[199] Keeve Brodman, 'Differential Diagnosis by Syndrome Complaints', Application for Research Grant (RG-4742), National Institutes of Health, 1 Feb. 1956, KBP/1/7.

and constrained by paper-based structures and practices that predated Brodman's interest in computerisation.

This is not to say, of course, that coding CMI data for processing by a digital electronic computer made no difference. After working for a period of time at Yale, van Woerkom left the academy and moved to the Electric Boat Company, a subsidiary of General Dynamics Corporation that manufactured submarines for the U.S. Navy. There, he served as a technical director and a primary manager of research into submarine combat systems. In these roles, van Woerkom 'introduced computers into engineering, scientific and business systems, and applications to the development of tactical weapons systems for submarines'.[200] He also had access to the company's IBM 704 computer, which proved useful in his collaboration with Brodman (Figure 5). Use of the IBM 704 assisted tremendously in making tabulations of massive amounts of medical statistical data. Once digitised, the CMI data could be more easily analysed and manipulated; the testing of different formulas for their mathematical model became far less labour intensive (which is not to say that significant labour was no longer required). In short, digitisation reduced the 'friction' of CMI data—it rendered a large quantity of patient data far more tractable.[201]

Yet, the fact that computerisation greatly facilitated engagement with CMI data does not diminish the observation that this data was derived from paper-based tools, originally existed on paper, and was first studied largely 'by hand analysis'. Nor did it mean that manual and paper-based work suddenly vanished. Even with access to

---

[200] Raynor Duncomb and Morris Davis, 'Obituary: A. J. J. van Woerkom, 1915-1991', *Bulletin of the American Astronomical Society* 23 (1991): 1495.
[201] On the role of digitisation in reducing data friction, see W. Patrick McCray, 'How Astronomers Digitized the Sky', *Technology and Culture* 55 (2014): 908-944; Sepkoski, 'The Database before the Computer?', *Osiris* 32 (2017): 175-201, pp. 178, 191.

automated modes of tabulation, and in spite of his grievances about the labours of manual calculation, Brodman occasionally ended up performing tabulations by hand. Sometimes this was out of necessity. Because Brodman, housebound by multiple sclerosis, did not have direct personal access to the IBM 704, if he wanted to quickly test an idea or formula that popped into his head, he would often perform preliminary calculations by hand. Other times, however, he turned to manual practices at his own discretion.

Brodman seemed to believe that manual data practices could engender a different kind of relationship between a researcher and his data than could computerized data practices. On at least one occasion, he recommended these paper data practices to his colleagues, suggesting that manual tabulation generated deeper reflection. In a 1958 letter to van Woerkom on the topic of testing their most recent model, Brodman wrote, 'As I told you over the phone this morning, I think it best if we tabulate the results by hand instead of machine, sin[c]e this will give us the opportunity to contemplate the data'.[202] This attitude about the importance of time ('opportunity') to contemplate data reveals how computers at this stage in the research project were largely performing calculations that Brodman and his colleagues could perform (and had already been performing) manually—only computers performed this work faster. Brodman felt that this speed, while most often an asset, also came with its drawbacks, robbing the researcher of time to engage with and reflect upon his data.

THE MEDICAL DATA SCREEN

The fruit of all this labour—a working statistical model suitable for computer implementation—was announced in a series of articles published between 1959 and

[202] Brodman to van Woerkom, 19 Jan. 1958, KBP/6/3.

1961.[203] These later articles described a four-step statistical method. First, using the 1948-1949 CMI data, it assigned each questionnaire item a significance value for each of the 60 most common diseases according to the following formula:

$$P_{ij} = \frac{R_{ij} - R_i}{2(\sqrt{R_i})} \pm 1,$$

where P_{ij} is the significance value, R_i is frequency of a given complaint i in the entire patient population, and R_{ij} is the frequency of the complaint in the subset of patients with disease j.[204] The logic behind this formula is that a given complaint (or symptom) is diagnostically valuable when its frequency in the overall patient population is very different from its frequency in the subset of patients with the disease in question. Second, for a particular patient, and in a given disease, the significance values of the separate questionnaire items are combined using the formula:

$$N_p = \sum {P_{ij}}^2.$$

Third, this total value, N_p, is corrected for the patient's age—a statistical nod to the fact that physicians expect diseases to present differently in differently aged populations. Finally, the corrected total score for the particular patient is converted into a 'likelihood score', L, by relating it to the mean corrected total score for all patients known to have the disease, N_m:

$$L = \frac{N_p}{N_m}\ x\ 50.$$

The patient in question is diagnosed with the disease if L surpasses a given threshold. Brodman and van Woerkom developed a programme to run these computations on an

[203] Keeve Brodman, Adrianus van Woerkom, Albert Erdmann, Jr., and Leo Goldstein, 'Interpretation of Symptoms with a Data-Processing Machine', *A.M.A. Archives of Internal Medicine* 103 (1959): 116-122.
[204] Keeve Brodman, 'Diagnostic Decisions by Machine', *IRE Transactions on Medical Electronics* ME-7 (1960): 216-219, p. 217. On statistical aspects, see Adrianus van Woerkom and Keeve Brodman, 'Statistics for a Diagnostic Model', *Biometrics* 17 (1961): 299-318.

IBM 704 computer, which was capable of making these computations at a rate of one patient per second.[205]

In presenting the model, Brodman speculated about the ways in which it might 'simulate' clinical cognition.[206] He and van Woerkom tested various formulas for calculating the total score (N_p) for each disease. Besides the formula that was eventually adopted—summing the *squares* of the significance values—one of the alternative formulas that they favourably considered involved simply adding together the significance values of all the pertinent symptoms. A number of these formulas seemed to work satisfactorily. But Brodman was drawn to the formula that squared the significance values because it captured what he believed to be an important aspect of clinical reasoning. Specifically, Brodman believed that physicians weigh symptoms that are hallmarks of different diseases especially strongly. According to Brodman, the formula simulated this clinical tradition by squaring the diagnostic value assigned to each item, which further inflated the diagnostic weight of items with large significance values (consider: $1^2 = 1$; $3^2 = 9$).[207] Brodman's engagement with the CMI, mathematical modelling, large quantities of data, and computers prompted him to generate new ideas about how doctors think. The MDS method would in turn come to contain—in its code—these very ideas. The computerisation of diagnosis was clearly a generative space for thinking about, and encoding, the medical mind and its inner working. As Lucy Suchman has noted, and as Stephanie Dick has reiterated in the context of mathematical proof, computers can be 'powerful disclosing agents'—

[205] Adrianus van Woerkom, 'Program for a Diagnostic Model', *IRE Transactions on Medical Electronics* ME-7 (1960): 220.
[206] Brodman, 'Diagnostic Decisions by Machine', p. 216.
[207] Ibid., p. 217.

prompting humans to look inward and articulate certain views of themselves.[208] Brodman's case illustrates that this is not just a process of disclosure. It is also one of inscription and propagation: views articulated by humans can come to be encoded in computer programmes and the mathematics underlying them.

Yet for all that Brodman said about the possible analogies between the model and the human mind, he was the first to admit that the model as published fell short of his earlier, roaming speculations. 'The machine cannot yet do as well as a physician', he admitted to van Woerkom in 1958. 'It doesn't have the background of a physician's information, and we don't know how to programme a method of analysis a human being uses'.[209] He struck similar notes of moderation in his article describing the model, noting that the human mind remained a 'black box', shut off from perfect analysis.[210] Due to setbacks, combined with the mounting pressure to produce something concrete from all this work, Brodman's focus increasingly moved away from the model's form—the extent to which it accurately mimicked human cognition—and toward its function—the extent to which it generated correct diagnoses.[211] This rift between the aspirational and the actual brings light to the path-dependent character of the MDS's development, whereby earlier paper-based practices and earlier stores of data constrained subsequent lines of work and research. Brodman's project of computerisation proceeded along tracks that had been laid before computers entered the picture.

---

[208] Lucy Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions* (Cambridge: Cambridge University Press, 2006), pp. 226-240; Dick, 'AfterMath'.
[209] Brodman to van Woerkom, 29 Mar. 1958, KBP/6/3.
[210] Brodman, 'Diagnostic Decisions by Machine', p. 219.
[211] This progression was not perfectly linear. For example, Brodman's musings about pattern recognition, discussed in the previous section, came after he and his colleagues published their first accounts of the 'machine method' in 1959. These discussions about pattern recognition, however, would not amount to anything concrete, and Brodman and van Woerkom soon returned to the machine method as originally published.

The 'machine method' as described between 1959 and 1961 served as the basis for what would become the MDS.[212] Only two features differentiated the MDS from these earliest descriptions of the method. First, the MDS screened for 100, rather than just 60, diseases. Second, the MDS made use of a new questionnaire: the Medical Data Index (MDI). Even as Brodman and his colleagues revised the original questionnaire for the purpose of the computerised method, the revised version retained the logic, structure, form, and many of the individual questions that characterised Brodman's earlier paper-based work. Like the CMI, the MDI favoured informal over technical language and grouped questions on the questionnaire form according to sections of the body. Although the MDI contained fewer questions than the CMI (Brodman omitted the questions that did not help differentiate among the 100 diseases), many of the MDI's 150 questions were taken from its ancestor questionnaire with minimal changes.[213] Most fundamentally, both the CMI and the MDI had paper bodies—both were paper-based questionnaires. The research team created an MDS Manual, analogous to the CMI Manual, that outlined the intended operation of the method. The patient would complete the self-administered, paper-based MDI questionnaire, the results of which would be mailed to those with access to a digital computer for analysis. After human workers coded and input the MDI data, the computer would generate a paper output, which would then be sent back to the physician. Finally—and most importantly in Brodman's eyes—the physician would interpret this output in light of his 'own training, knowledge, experiences and insights' (Figure 6).[214] The MDS

[212] Keeve Brodman and Adrianus van Woerkom, 'Computer-Aided Diagnostic Screening for 100 Common Diseases', *JAMA* 197 (1966): 179-183.
[213] Brodman to Barbara Bates, 18 Oct. 1967, KBP/4/2.
[214] Keeve Brodman, *Manual: Medical Data Screen Index (MEDIS)* (Medical Data Corporation, 1964), p. 11, KBP/1/11.

method, as originally conceived, therefore continued to be heavily reliant on paper technologies—and on the circulation of paper among various actors.

By the mid-1960s, Brodman had 'become convinced … that there is an urgent need to have our method in operation and for sale as quickly as possible'.[215] To help bring the MDS into wider clinical use, Brodman created the Medical Data Corporation, a small company staffed by Brodman, his research associates, and a few additional personnel to handle business and administrative affairs. Though established in 1961, the Medical Data Corporation ramped up its activities in 1966, after the MDS method was finalised and reported in the medical literature. The Medical Data Corporation focused mostly on bringing the MDS to large companies. This focus on industrial health was somewhat circumstantial, driven by Brodman's personal interest and expertise in the field. But Brodman and his associates also had good reason to believe that industry represented a particularly suitable area of application. Large companies had strong incentives to maintain and monitor the health of their employees, and they often lacked well-developed or well-equipped medical departments to do so. The Medical Data Corporation initiated discussions about the MDS method with a handful of large corporations, many with tens of thousands of employees: General Electric, General Motors, Olin Mathieson, Mobile Oil, Owens-Illinois, New York Telephone Company, Prudential Life Insurance Company, Mountain States Telephone Company, and Tennessee Eastman Company—among others. As the Medical Data Corporation described it, companies would use the MDS method as follows. First, the Medical Data Corporation would provide the company with the needed quantity of MDI questionnaires at a cost of $4.10 per 50 questionnaires. After having employees complete the MDI, companies would send the completed questionnaires to the Medical

[215] Brodman to van Woerkom, 30 Jul. 1964, KBP/6/4.

Data Corporation for servicing. The Medical Data Corporation would then use a computer programme to screen the questionnaires for the 100 common diseases, sending the resultant computer output back to the company at a cost of $1.00 per report.

Promoting the MDS gave Brodman and his associates a good sense of how different audiences viewed their effort. The reactions provoked by the MDS echoed those provoked by the CMI some two decades earlier—only now they were amplified. This intensified response resulted from two key differences between the MDS and the CMI. First, the MDS made use of a technology, the computer, about which many physicians and lay people had negative or uneasy preconceptions.[216] Second, the MDS method, by its very nature, contradicted Brodman's earlier assurances in the CMI manual about the physician's interpretive monopoly over medical data. In the late 1940s, Brodman was reassuring the medical community that the interpretation of symptoms belonged to the physician and the physician alone; some ten years later he was developing 'a machine method of making diagnostic interpretations' and publishing articles with titles like 'Interpretation of Symptoms with a Data-Processing Machine'.[217]

Pushback came as soon as Brodman and his collaborators began seeking out research funding. In 1962, after hearing rumours about the availability of 'substantial sums of money to spend for medical research', Brodman wrote to a research branch of the Army Medical Research and Development Command to inquire about funding for

---

[216] On lay attitudes towards computers, see, e.g., Sarah Igo, *The Known Citizen: A History of Privacy in Modern America* (Cambridge, MA: Harvard University Press, 2018), pp. 232-263.

[217] 'Operation Diagnostic Screen', 'Rapid screening of VA patients for 100 diseases', draft application, n.d. (proposed project dates: 1 May 1961 to 30 Apr. 1963), KBP/3/12.

his work.[218] The response to Brodman's inquiry was polite but blunt: '[S]ome of us are rather vigorously opposed to this sort of thing. I am sure you must have experienced this yourself from other members of the medical profession'.[219] Brodman and his colleagues faced similar opposition once they started submitting their work to medical journals. In 1960, an anonymous reviewer for *Biometrics* tentatively recommended Brodman and van Woerkom's article for publication, in part as a way of recognising 'the courage to undertake such a difficulty study' and in part as a way of 'stimulat[ing] interest and controversy'. But this recommendation came with a caveat: 'I tremble at the thought that someone might be persuaded to use this procedure for practical purposes'.[220]

It did not take long for Brodman to appreciate what he called 'the expressed and unconscious antagonism' of many physicians to a machine method of diagnosis.[221] But a deeper understanding of the specific concerns underlying this antagonism came primarily after the Medical Data Corporation started exploring the potential market for the MDS. In general, the Medical Data Corporation found that the '[d]octor feels [that the] program may jeopardize his position or job'. The MDS, industrial physicians believed, 'offers a threat'.[222] If the MDS didn't take their job, many physicians thought the method risked changing it in detrimental ways. The medical director of Mobile Oil related to the Medical Data Corporation the opinion of one company physician 'who

---

[218] Max M. Kampelman to Brodman, 13 Feb. 1962, KBP/3/12.
[219] Philip Sperling to Brodman, 26 Feb. 1962, KBP/3/12.
[220] Ralph Bradley to Brodman, 11 Nov. 1960, KBP/3/12.
[221] Brodman to Kampelman, 12 Jan. 1960, KBP/3/12. While the name Max Kampelman is written at the top of this letter, its author and recipient are not definitively clear. From its content and style—and its context within other letters sent between Brodman and Kampelman—I have inferred that it is a letter from Brodman to Kampelman.
[222] Arnold Wander to Herbert Krasnow, 5 Dec. 1967, KBP/4/2. This letter was sent to Krasnow via his secretary. See also Wander to Krasnow, 22 Sep. 1967, KBP/4/2.

felt that the computer would replace his personal contact with the patient'.[223] A magazine article on the MDS method gave voice to this same concern, reporting the fear among some physicians 'that increased mechanization of the diagnostic phase of medical practice will dehumanize the doctor-patient relationship'.[224]

General practitioners at Owens-Illinois, a large glass company, bristled at the thought of the MDS—a reaction that the company's medical director attributed to the fact that general practitioners were 'not statistically minded' enough 'to accept the MDS as a feasible approach'.[225] While Brodman believed that much of the opposition to the method stemmed from fear and ignorance, other critiques of his method were in fact quite practical. One repeatedly-voiced concern was that the time lag between (1) sending the questionnaire responses to the Medical Data Corporation for computer analysis and (2) receiving the results of that analysis would be too great for the service to be effective and practicable. There was also a healthy degree of scepticism around the claim that the MDS would save companies, hospitals, and physicians large sums of money. The cost of the MDS, some industrial physicians noted, was not just the cost of leasing the programme itself. They also worried privately about the added cost of having to pursue medical problems that had otherwise gone unrecognised. Critics today might draw comparisons to the cost of pursuing one 'incidentaloma' after another.[226] Failing to follow up on the many problems identified by the MDS, moreover, might put

---

[223] Brodman to Wander, 27 Sep. 1967, KBP/4/2.
[224] 'Computers Programmed to Sort Routine Symptoms', n.d., The Ralph Engle, MD (1920-2000) Papers, box 8, folder 2, Medical Center Archives of New York-Presbyterian/Weill Cornell, New York (hereafter cited in the form REP/8/2).
[225] Wander to Krasnow, 30 Nov. 1967, KBP/4/2.
[226] G. W. Geelhoed and E. M. Druy, 'Management of the Adrenal "Incidentaloma"', *Surgery* 92 (1982): 866-874. The term 'incidentaloma' was coined in the 1980s, as computer tomography (CT) scans were becoming increasingly available. The term refers to incidental findings for which there exists no corresponding symptoms. It sheds light on the complexities around overdiagnosis and overtreatment.

the physician in legal jeopardy. This view was expressed in March of 1968, during a meeting between MDS sales manager Arnold Wander and a half-dozen industry physicians. As Wander reported, 'Confidentially their attitude was, why look for trouble?'[227] It was possible, these physicians implied, for the MDS to do its job too well.

As with the CMI, Brodman included in various MDS publications assurances intended to defang the MDS method as it existed in many physicians' minds. For example, Brodman laboured over the wording of one paragraph in the MDS Manual, a paragraph he included for the sole purpose of appeasing uneasy physicians. As Brodman explained in a 1972 letter, 'I want this paragraph to be as clear as possible because it is being entered to allay the fears of those physicians who may feel that an automated method of identifying symptom complexes threatens to replace them'.[228] This paragraph cast the computer as subservient to the physician, stressing that by using the MDS physicians were merely 'delegat[ing] to a computer the rote task of matching symptoms while reserving for himself the clinical evaluation of data which the MDS develops'. The passage also enumerated the computer's limitations, a device 'totally incapable' of making sense out of uncertain or inconsistent data.

Whatever their effect on computerphobes, computerphiles saw straight through such assurances. In a 1966 manuscript submitted to *JAMA*, Brodman and his coauthors assured that computers would not partake in value judgements anytime in the near future. One of the manuscript's anonymous reviewers found this statement indefensible: 'The comments … to the effect that value judgment will not be made by the computer in the forseeable [*sic*] future will be balm to the ears of many an M.D. in

[227] Wander to Krasnow, 30 Nov. 1967, KBP/4/2.
[228] Brodman to Harold Dehner, 28 Apr. 1972, KBP/5/1.

the U.S. The fact is that such an opinion can only be held (1) for the purpose of soothing the obvious fears, (2) out of ignorance, or (3) by an intellectual Luddite. Clearly, in the future, the computer will in many cases be making better value judgments than the average M.D'.[229] This reviewer's response brings out the degree of variability in attitudes toward the MDS. After Brodman published his first article on the MDS method in 1966, letters poured in from around the country with reactions that ranged from curious intrigue to outright enthusiasm. Some physicians happily imbibed Brodman's vision of the MDS as an invaluable adjunct to the busy physician. Some countered that technologies like the MDS would, in fact, replace physicians but reasoned that such an outcome wasn't necessarily bad. The steady advance of technology would not be kind to the sentimental, they warned. Some hoped that the MDS might one day stand in for doctors but only in 'underdeveloped' areas without an adequate supply of medical personnel.[230] Physicians committed to various medical projects—telemedicine, regional medicine, preventative medicine—also envisioned a productive place of the MDS in their respective undertakings.

The MDS method appealed to many companies as well. But despite some nibbles, no one bit. The sticking point was turning this general interest into an actual commitment to adopt the tool. By late 1967, the Medical Data Corporation still had not made a sale. What bothered Brodman and his colleagues most about this situation was the squandered potential. At least one Medical Data Corporation associate saw the MDS as 'a potential million dollar enterprise'.[231] Sales manager Arnold Wander captured this general feeling nicely in a letter to another company official: '[I]t is not fair to this exciting program to simply "spin ones wheels" [*sic*]'. Though frustrated and

[229] Brodman to Charles Roland, 23 May 1966, KBP/4/1.
[230] Robert Albrecht to Brodman, 5 Dec. 1968, KBP/4/3.
[231] Wander to Krasnow, 22 Jan. 1969, KBP/4/3.

disappointed, Wander reiterated his commitment to the MDS and its mission: 'I am not disillusioned; have not lost my enthusiasm nor my resolve to make a sale and get [the] MDS soaring'.[232] Success evaded the Medical Data Corporation for manifold reasons; among them were poor coordination among Brodman's associates, reluctance among some industry physicians, the company's limited capital, and lingering imperfections in the operation of the MDS method. Moreover, the Medical Data Corporation found itself in what seemed like an impassable catch 22: many industry leaders expressed interest in the MDS, but no one wanted to risk adopting it first. Invariably, these officials would ask Brodman and his associates to direct them to examples of other similarly situated companies that had successfully implemented the MDS.

In spite of the Medical Data Corporation's troubled efforts to sell and distribute the MDS, the computerised method enjoyed some positive media coverage. The *Medical World News*, a magazine for the medical profession with a circulation of 250,000, published an article on the MDS in its 16 February 1968 issue.[233] The article opened in a distinctively journalistic vein: 'The electronic brain is now being readied for a move into the vital core of a doctor's practice—the diagnosis of his patients' disease'.[234] Such journalistic flourishes aside, the article, based largely on interview material, provided Brodman with a platform to communicate his views on the MDS to a wider medical audience. From Brodman's perspective, the article was imperfect, containing 'only very sketchy [partial] information about the method'.[235] But the Medical Data Corporation was nevertheless pleased with the coverage, with one

[232] Wander to Krasnow, 6 Nov. 1967, KBP/4/2.
[233] In 1965, the *Medical World News* had estimated circulation between 205,000 and 230,000. 'Challenging the Leader', *Time* 85, no. 5 (1965): 71.
[234] 'Computers Programmed to Sort Routine Symptoms', n.d., REP/8/2.
[235] Brodman to Richard Schwartz, 8 Mar. 1968, KBP/4/3.

associate reasoning, 'Since this publication has a circulation of nearly a quarter of a million, this exposure is extremely valuable…'.[236]

More good news came a few months later, when Brodman received a letter from Marvin L. Miller, the director of Roche Psychiatric Service Institute, an arm of the multinational healthcare company Hoffmann-La Roche. Hoffmann-La Roche had recently been involved in creating and distributing an automated version of the MMPI, and Miller indicated that the company might be interested in doing something similar with MDS. Managers at the company, Miller noted, were 'extremely interested in exploring the possibility of making such an automated service available nationally'.[237] Brodman jumped at this opportunity—and understandably so. In light on the Medical Data Corporation's limited, disorganised, and largely ineffectual efforts to distribute the MDS, the prospect of a large healthcare company orchestrating a national rollout of the programme surely appealed to Brodman and his colleagues. By late spring, representatives of the two organisations had hammered out a formal agreement; the Medical Data Corporation terminated its independent efforts to lease the MDS and granted Hoffmann-La Roche exclusive rights to pilot test the program.

A PRONG OF THE 'MANY-SIDED ATTACK'

What value did Miller, Brodman, and their colleagues see in the MDS? And why? What led Brodman and his people to conclude, in the mid-1960s, that there was 'an urgent need to have our method in operation and for sale as quickly as possible'? Optimism about the MDS was buoyed by a number of medical reform efforts underway during this time period. On 10 February 1964, President Lyndon Johnson delivered his

[236] Joseph A. Louis to Wander, 22 Mar. 1968, KBP/4/3.
[237] Marvin L. Miller to Brodman, 26 Apr. 1968, KBP/4/3.

'Special Message to the Congress on the Nation's Health', which set forth his Administration's health care agenda for the coming year. The address drew wide popular attention primarily for what it said with respect to the establishment of a national insurance programme for aged Americans—Medicare. Although Medicare may have occupied the top of the President's agenda, Johnson also packed into the address a number of other significant policy proposals that together were to 'comprise a vigorous and many-sided attack on our most serious health problems'.[238]

Among the largest targets of Johnson's 'many-sided attack' were three diseases: heart disease, cancer, and stroke. These diseases, Johnson pointed out, afflicted 15 million Americans and had long remained the leading causes of death in the United States. Johnson believed that the health challenges posed by these diseases, and their possible policy remedies, demanded further study. To this end, Johnson announced the creation of a Commission on Heart Disease, Cancer, and Stroke 'to recommend steps to reduce the incidence of these diseases through new knowledge and more complete utilization of the medical knowledge we already have'.[239] A few months after giving his health message to Congress, Johnson tapped Michael DeBakey, an innovative and internationally renowned cardiac surgeon at the Baylor College of Medicine, to head the Commission. In December 1964, the twenty-eight-member Commission published its findings and recommendations in an expansive two-volume report, A National Program to Conquer Heart Disease, Cancer and Stroke, which Johnson received at a ceremony in the Cabinet Room of the White House. News outlets recognised the Report's significance and magnitude; the day after its publication, the *Washington Post* ran the front-page headline, 'Experts Ask U.S. War on 3 Diseases'.

---

[238] Lyndon B. Johnson, 'Special Message to the Congress on the Nation's Health', 10 Feb. 1964.
[239] Ibid.

Four recommendations made up the substantive backbone of the DeBakey Report. Collectively these recommendations—the establishment of (1) 'Regional Heart Disease, Cancer and Stroke Centers … in universities, hospitals and research institutes', (2) 'Diagnostic and Treatment Stations in communities across the nation', (3) 'medical complexes [stimulated via grant support] whereby university medical schools, hospitals and other health care and research agencies and institutions work in concert', and (4) 'developmental grants to medical schools to enable these institutions to improve their total capability for both academic and research programs'—set up a coordinated system of resources designed to function cooperatively and hierarchically.[240] Coordination among these different resources, as envisioned by the Report, looked something like this: 60 Regional Heart Disease, Cancer and Stroke Centers (1) would, in combination with grant-supported medical schools (4), assemble multidisciplinary research and clinical teams to perform clinical investigations, conduct fundamental research, and (to a lesser extent) provide advanced and complex patient care. The new methods of diagnosis and treatment generated by these Centers would disseminate out to regional communities through 450 Diagnostic and Treatment Stations scattered around the country (2). With federal grant support, pre-existing nuclei of medical research (3) would be expanded and integrated into this hierarchical network of Centers and Stations. The Regional Medical Program (RMP) was born.

In proposing this kind of institutional regime, the report espoused a logic that had guided health policy and reform throughout the twentieth century. This logic was one of 'hierarchical regionalism', whereby the knowledge, practices, and technologies developed in academic laboratories and teaching hospitals would spread down of

---

[240] The President's Commission on Heart Disease, Cancer and Stroke, *Report to the President: A Program to Conquer Heart Disease, Cancer and Stroke*, vol. 1 (Washington, D.C.: U.S. Government Printing Office, 1964), pp. 29, 32, 34-35.

hierarchies of schools, hospitals, clinics, and private practices and out to regional locales.[241] As the DeBakey Report framed it, the proposed national network would 'link every private doctor and every community hospital to a national—and indeed worldwide—network transmitting the newest and best in health service'.[242] Layered over this integrated system of medical sites and resources were an additional thirty-one recommendations. These additional recommendations touched on everything from incentive grants and public awareness campaigns, to biomedical research and continuing medical education. The report's drafters felt that the enormity of the problem posed by heart disease, cancer, stroke—diseases that, taken together, accounted for 'more than 7 out of every 10 deaths' in the United States—demanded an expansive and wide-ranging set of solutions.[243]

Never in the history of American medicine has a major healthcare reform effort escaped significant pushback, and the RMP was no exception to this pattern.[244] In fact, for many physicians, the RMP seemed to pose a professional threat even greater than that posed by Medicare: whereas the latter programme reconfigured the *financing* of health care, the RMP intervened directly in the *delivery* of care. As Assistant Secretary of HEW Wilbur Cohen indicated in a 1965 memo, the RMP, owing to its mission of 'completely reorganiz[ing] medical care in the United States', was seen by many within

---

[241] Daniel M. Fox, *Health Policies, Health Politics: The British and American Experience, 1911-1965* (Princeton, NJ: Princeton University Press, 1986).
[242] President's Commission, *A Program to Conquer*, pp. 28-29.
[243] Ibid., p. xi.
[244] On the history of the RMP, see Stephen P. Strickland, *The History of Regional Medical Programs: The Life and Death of a Small Initiative of the Great Society* (Lanham, MD: University Press of America, 2000). On the RMP in the context of Johnson's other healthcare reform efforts, see Paul Starr, *The Social Transformation of American Medicine: The Rise of a Sovereign Profession and the Making of a Vast Industry* (New York: Basic Books, 1982), p. 370.

the medical community as 'more radical' than Medicare legislation.[245] These hostile attitudes were expressed most vigorously by the American Medical Association (AMA), an organisation that controlled many levers of influence over the legislative process. During the 1960s, the AMA and its surrogates testified before Congress; liaised privately with members of both the executive and the legislative branches; and supported, via the American Medical Political Action Committee (AMPAC), the political campaigns of favoured congressional candidates. Through these and other levers, the AMA worked to translate its disapproval into substantive changes to the proposed legislation—if not to kill the legislation altogether. While the AMA failed to stop the RMP legislation, the group effectively tempered many of its reforms. The bill that Johnson signed into law in many ways represented only a bare tendril of the programme originally envisioned by the DeBakey Commission. At each subsequent step in the legislative process, the programme atrophied: the DeBakey Commission conceived of a five-year, $2.94 billion program; the Johnson Administration recommended a five-year, $1.15 billion program; the Senate bill proposed a four-year, $650 million program; and, finally, the House bill authorised—and Johnson ultimately signed into law—a three-year, $340 million programme.

Though watered down, the programme was not trivial. Indeed, the RMP funded a host of early projects in the computerisation of medicine and medical diagnosis. The RMP's support for such efforts resulted from the understanding that computerised systems could advance the program's larger 'circulatory' aims: computerised systems, many believed, might help move the knowledge and expertise generated at the 'heart' of the medical system (medical schools and teaching hospitals) out to the system's

[245] Quoted in Robert Dallek, *Flawed Giant: Lyndon Johnson and His Times, 1961-1973* (Oxford and New York: Oxford University Press, 1998), p. 675.

regional ‘extremities’ (patients and clinics in rural America). Among the more colourful of such RMP-funded initiatives was the Automated Physician’s Assistant (APA). Designed to ‘bring modern medicine out to the rural areas in Missouri’, the APA was a pilot programme in the town of Salem, Missouri (population 4,363), whereby the office of the local, country physician, Billy Jack Bass, was equipped with automated medical history taking systems and other information acquisition devices.[246]

Although the MDS did not receive any RMP funds directly, many viewed the method within the context of these larger efforts. Brodman surely did: he urged Roche personnel to ‘consider the opportunities in other important areas as, for example … in government-supported efforts, such as the Regional Medical Program or Welfare Program’.[247] Brodman’s interest in the RMP was reciprocated by RMP personnel’s interest in the MDS: certain members of various RMP initiatives looked on the computerised method with great interest. In 1968, a leader of the Missouri Regional Medical Program wrote to Brodman, asking about the MDS and its possible incorporation with their efforts in computerised medicine: ‘I have just discovered your article on the Medical Data Screen and am very interested in it. Here in the Missouri Regional Medical Program, we are working on an Automated Patient History Acquisition System (project APHAS) [part of the APA project discussed above] and some of the data you allude to would be very helpful in designing our program’.[248]

---

[246] On Billy Jack Bass and the Automated Physician’s Assistant, see Arthur E. Rikli, Fred V. Lucas, and Fred Frazier, ‘The Automated Physician’s Assistant’, *Clinical Engineering News* 3, no. 4 (1975): 3-5; John G. Rogers, ‘Dr. Billy Jack Bass: He’s “Helping to Make Medical History”’, *The Palm Beach Post* (West Palm Beach, FL), 11 Jul. 1971, p. 138. See also Interview with Dr. Arthur E. Rikli, 25 Jul. 1991, The Regional Medical Programs Collection, U.S. National Library of Medicine, National Institutes of Health, Bethesda, Maryland. Transcript available online at: https://profiles.nlm.nih.gov/ps/retrieve/ResourceText/rmggac.

[247] Brodman to Miller, 30 Apr. 1968, KBP/4/3.

[248] Gerald Fleischli to Brodman, 3 Oct. 1968, KBP/4/3.

Such RMP initiatives were but a single instantiation of a broader ethos. Owing to the rapid pace of technological change, many began to imagine the computer as an antidote to many of medicine's steepest challenges. Such perspectives were voiced prominently in the 1966 Report of the National Commission on Technology, Automation, and Economy Progress. Travelling under the title 'Technology and the American Economy', the congressionally commissioned report aimed to define areas of 'unmet community and human needs' and recommend ways of deploying new technologies to meet said needs.[249] Against the backdrop of such national discussions, an anonymous reviewer of one of Brodman's manuscripts picked up on the MDS's broader timeliness: 'I find this paper informative, interesting, and timely—particularly in the light of (1) the increasing strain that will shortly be placed on medical manpower by Medicare and the Regional Medical Programs, (2) the Adult Health Protection Act of 1966 currently being proposed by Congressman Fogarty and Senator Williams, and (3) the burgeoning interest in the possibility of Regional Health Computer Systems as indicated in the February 1966 report of the President's Commission on Technology, Automation, and Economic Progress'.[250] Indeed, the report identified four main recommendations to improve health care in the United States, two of which were very closely aligned with the CMI and the MDS: (1) 'broader and bolder use of the computer and other new health technologies' and (2) 'increased spread and use of health statistics, information, and indexes'.[251]

---

[249] National Commission on Technology, Automation and Economic Progress, *Technology and the American Economy*, vol. 1 (Washington, D.C.: U.S. Government Printing Office, 1966), p. xiv.
[250] Brodman to Charles Roland, 23 May 1966, KBP/4/1.
[251] National Commission on Technology, Automation and Economic Progress, *Technology and the American Economy*, Vol. 1 (Washington, D.C.: U.S. Government Printing Office, 1966), pp. 78-83, 112.

Besides the MDS's alignment with ongoing national reform efforts, Brodman viewed the computerised method—and its paper ancestor—as well positioned to solve some of the most pressing medical problems of the time. For Brodman, the primary trouble with the medical system was a lack of resources—human resources in particular. There were too few doctors with too little time. According to Brodman, the numerical imbalance between the provider and consumer sides of healthcare manifested in the form of rushed doctors and hurried clinics. Under the pressure of time, physicians tended to examine patients narrowly, taking into account only their most immediate complaints and problems. Any foray a physician hoped to make into additional areas of health risked further draining already scarce resources. As Brodman explained to one correspondent in 1969, 'Heavy patient loads prevent physicians or clinics from obtaining and interpreting comprehensive medical history on any but a few patients. Physicians thus must treat patients for their chief complaints and without reference to their total medical problems'.[252]

This concern about a national shortage of medical personnel was shared by many during the 1950s and 1960s. The country's perceived shortfall of doctors was a dominant theme of the 1952 report by the President's Commission on the Health Needs of the Nation. And the issue surged again onto the national agenda a few years later, in 1959, after the publication of a report by the Surgeon General's Consultant Group on Medical Education. Named after the consultant group's chairman, Frank Bane, the 'Bane report' projected that the country would face a deficit of 40,000 doctors by 1975.[253] The report attributed medicine's manpower problem to two principal causes.

---

[252] Brodman to Charles Burger, 24 Sep. 1969, KBP/4/4.
[253] *Physicians for a Growing America*, Report of the Surgeon General's Group on Medical Education, publication No. 709 (Bethesda, MD: U.S. Department of Health, Education, and Welfare, 1959).

The first was numerical: there were simply too few doctors in the field and too few trainees in the pipeline. This numerical problem was further compounded by a distributive problem: on top of the fact that there were too few doctors overall, the distribution of this limited supply of personnel was not optimised to meet the nation's health needs. Unacceptably large shares of physicians pursued some kind of specialty training over general practice.

To some degree, the Bane report achieved its objectives, spurring a degree of medical educational reform not seen since the Flexnor report.[254] The Bane report's findings also made waves through the popular press. In 1960, *Harper's Magazine* ran a 42-page special supplement on 'The Crisis in American Medicine', parts of which addressed the shortfall of doctors as outlined in the Bane report: 'Shrinking fastest of all is the supply of personal physicians'. Framing the problem more numerically, the article continued, 'If internists were to provide continuous care for all patients there should be at least one internist for every 750 persons. Actually, the ratio is less than one per 14,000, and they are not evenly distributed throughout the country'.[255]

Brodman fashioned the CMI and the MDS as partial solutions to this problem. Neither the CMI nor the MDS deployed any more boots on the ground, of course. But they did aim to assist, extend, and optimise the medical manpower that was already in place. These tools allowed the busy physician to gain a wider perspective on the patient. And, perhaps as importantly, they offered this perspective 'with no expenditure of the short supply of physicians' time'.[256] 'The Medical Data Screen', Brodman

---

[254] Kenneth Ludmerer, *Time to Heal: American Medical Education from the Turn of the Century to the Era of Managed Care* (Oxford and New York: Oxford University Press, 1999), pp. 210-212.
[255] Selig Greenberg, 'The Decline of the Healing Art', *Harper's Magazine* 221 (1960): 132-150, p. 147.
[256] Staff of the Medata Foundation and Medical Data Corporation, unpublished notes, 1 Jun. 1962, KBP/3/12.

elaborated in 1967, 'is unique in bringing comprehensive information to the physician about a person's total medical status, and in doing so rapidly, inexpensively, and with little burden to physician or patient'.[257] The patient as revealed by the CMI and the MDS was not the 'diabetic patient' or the 'cancer patient' but, as Brodman put it, 'the total patient'.[258] Elsewhere, Brodman broadened 'the total patient' even further—to 'the whole person'.[259] Much has been said about the tendency of modern medicine to reduce the patient into his disease or collection of medical signs and symptoms. Brodman hinted that the CMI and the MDS corrected this tendency, redirecting the physician's attention from 'a particular symptom' to 'the total patient', perhaps even 'the whole person'. In the frantic, and increasingly compartmentalised, world of medicine, the MDS conferred both speed *and* completeness.

As strongly as Brodman emphasised the importance of seeing 'the total patient', he characterised his initiatives—particularly the MDS—as a species of laboratory test. The MDS 'is essentially a laboratory technique', he noted frequently.[260] Extending this analogy, he suggested in 1968 that sending the patient's MDI responses to a computer for analysis was roughly equivalent to sending the patient's blood chemistry to a laboratory for analysis.[261] Brodman pointed out, however, the MDS was a laboratory test of a distinctive kind. Existing laboratory tests were typically limited in scope. To obtain even a reasonably comprehensive picture of the patient, the physician had to

---

[257] Brodman to Marvin Sameth, 6 Oct. 1967, KBP/4/2.
[258] 'Machine diagnostic screening of 10,000 patients', Application for Research Grant (RG-9116), National Institutes of Health, 17 Mar. 1961, KBP/1/7.
[259] Brodman to Sameth, 6 Oct. 1967, KBP/4/2.
[260] 'Adjunct for comprehensive medical diagnosis', Application for Research Grant (No. GM 14382-01), Public Health Service, 1 Feb. 1966, KBP/1/7.
[261] Keeve Brodman, Minutes of Meeting with Dr. Kenneth T. Bird, 30 Oct. 1968, KBP/4/3. On Kenneth Bird and his initiatives in telemedicine, see Jeremy A. Greene, 'When Television Was A Medical Device', *Humanities: The Magazine for the National Endowment for the Humanities* 38, no. 2 (2017).

order multiple costly tests.[262] Brodman and his associates maintained that the MDS 'avoid[ed] [this] "shotgun" approach', garnering information on multiple diverse common diseases and conditions with a single 'test'.[263]

By claiming that the MDS was a laboratory test of a unique sort, Brodman was communicating something important: the MDS is at once exactly like *and* totally unlike anything physicians had seen before. In terms of the physician's daily practice, the MDS would not change a thing. But in terms of the information that it furnished, Brodman signalled that the instrument's impact would be profound. In other words, framing the MDS as a laboratory test conveyed that this new tool would improve medical decisions and diagnoses by developing upon—rather than disrupting—conventional modes of medical practice. More than anything, the likening of the MDS to a laboratory test represented the familiar strategy to allay the fears of wary physicians. In fact, the Medical Data Corporation's sales manager confessed as much in a 1967 correspondence: 'Dr. Brodman suggests that the advertising program … be alerted so that they can emphasize that our MDS is a program of laboratory technique designed to alert the physician to disease present and in no way diagnosis [*sic*] or treats and in no way endangers doctors position as the final factor in the determination of disease present or treatment'.[264]

Comparisons between the MDS and laboratory tests thus provided arrows in the quiver as Brodman and his colleagues defended their computerised method against common critiques. In August 1968, the head of the Regional Medical Program in Texas voiced an objection to the MDS that was later characterised during a conference

---

[262] 'Adjunct for comprehensive medical diagnosis', Application for Research Grant (No. GM 14382-01), Public Health Service, 1 Feb. 1966, KBP/1/7.
[263] Wander to Krasnow, 7 Nov. 1967, KBP/4/2.
[264] Wander to Krasnow, 5 Dec. 1967, KBP/4/2.

between Roche and Brodman's team: 'He said that the one objection he has to the Medical Data Screen is that it may lead physicians to give up the practice of thinking and let the computer make diagnoses'.[265] The computerisation of medicine, some worried, might slowly blunt physicians' clinical skill and acumen. Tools like these might be used as an aide, but they could also be used as a crutch.

Brodman and his colleagues countered this line of thinking by invoking the now-accepted reliance on laboratory tests. The mainstream medical community had long accepted laboratory tests as central components of cares—occasional bouts of overreliance notwithstanding. Why should the computer, Brodman and his colleagues wondered, be held to a different standard? 'This comment', they wrote, 'could be made about any effective laboratory. Indeed, physicians now too often rely on laboratory tests instead of making an extensive investigation of each patient and the interpreting the data collected. … We have always emphasized that the MDS is no more than a laboratory test and should be treated as such even though the results are reported in terms of named disease'.[266]

THE STATISTICAL PATIENT

Challenges to the CMI and the MDS, as described previously, sometimes rested on the humanistic dimensions of medicine. To some, both the CMI and the MDS harbingered a deadening of the clinical encounter—a reduction of the patient into binary variables and a distancing of the doctor from his patients. Yet Brodman, having been trained in psychiatry and psychosomatic medicine, was not only sensitive to the goals of holistic medicine but convinced that the MDS was consistent with them. If one

[265] Notes on Conference with Dr. Koriter of Roche, 16 Aug. 1968, KBP/4/3.
[266] Notes on Conference with Dr. Koriter of Roche, 16 Aug. 1968, KBP/4/3.

facet of the total patient was particularly overlooked by modern medicine, it was the emotional facet. In Brodman's eyes, this oversight was significant, for he 'strongly suspect[ed]' that patients' suffering and illness lies 'in large measure with the emotions as well as with organic disorders'.[267] He found the medical systems' uneven response to different kinds of illness and suffering troubling. '[I]t is sadly true', Brodman wrote in 1965, 'that whereas a person in our civilized society can usually get help for his bodily disorders he can rarely find help for disorders of the spirit and feelings'. Brodman evidently felt that the CMI and the MDS, with their capacity to elucidate patients' psychiatric and emotional status, could bring such persons into the medical fold. These instruments, moreover, liberated physicians from the monotonous task of conducting a full review of systems, freeing up time to engage with the patient as an individual. These sentiments demonstrate that even as Brodman blazed forward on projects of computerisation and standardisation, he continued to hold many ideas about medicine that he had developed earlier, through his training in psychiatry and psychosomatics. Indeed, in notes from 1971 regarding possible revisions to the MDS Manual, he reiterated his early view that 'many pts [patients] need help for emots [*sic*], often sufficiently supplied by MD or aide just through encouraging pt to talk and verbalize his troubles'.[268]

Others who were enthusiastic about bringing the computer into medical diagnosis and history taking shared similar attitudes. One of Brodman's correspondents, the (rather eccentric) Dutchman H. J. van Heyst, spoke repeatedly about his clinical interest in the aspects of human illness that went unrecognised by modern medicine—'the hidden segment of the "iceberg" of human suffering', in van

[267] Brodman to van Heyst, 23 Feb. 1965, KBP/3/13.
[268] Brodman, handwritten notes on MDS output, 30 Nov. 1971, KBP/2/10.

Heyst's words.[269] To uncover this large share of human suffering that remained under the surface, van Heyst counterintuitively turned to instruments of standardisation and automation. As he related to Brodman, 'Ultimately only a combined effort, using to maximum extent electronic data processing and standardization of procedure, will it become possible to deal with this "iceberg."'[270] This phrase—the iceberg of human suffering—was one that van Heyst borrowed from the famous Midtown Manhattan Study, a pioneering psychiatric and sociological study that aimed to unpack the relationship between urban health and the social community.[271] By van Heyst's estimation, the Midtown Manhattan Study was among the first and the best exemplars of 'comprehensive or holistic medicine'.[272]

As much as Brodman and others contemplated medical holism or the individuality of each patient, the MDS method in fact subordinated the eccentric, the individual, and the particular to the aggregate, the average, and the statistical. Brodman believed that the MDS's object of study was 'the total patient'. A better description of its object of study, however, would be 'the statistical patient': the MDS method generated statistical characterisations of patients and the disease categories to which they belonged. Brodman described these statistical renderings of diseases as symptom complexes. He came to the idea that certain groupings of symptoms were real and diagnostically meaningful entities early on in his research. At this early stage, Brodman's language in describing these groupings was fluid. He discussed 'clusters' or 'constellations' of complaints alongside 'syndromes' of symptoms. But by the late 1960s, Brodman was consistently using the term 'symptom complex' to describe the

[269] van Heyst to Brodman, 19 Feb. 1965, KBP/3/13
[270] van Heyst to Brodman, 11 Mar. 1965, KBP/3/13.
[271] Leo Srole et al., *Mental Health in the Metropolis: The Midtown Manhattan Study* (New York: McGraw-Hill, 1962).
[272] van Heyst to Brodman, 2 Sept 1965, KBP/3/13.

groupings of questionnaire complaints found in specific diseases. For each disease, Brodman could list its associated complex of symptoms. For instance, the symptom complex for disease F572 (chronic enteritis or ulcerative colitis) with respect to male patients included three items with a significance value of three—047 (intestinal trouble), 051 (bloody bowel movements), and 052 (pain on moving the bowels)—three items with a significance value of two—044 (indigestion), 049 (diarrhea), and 050 (mucus in stools)—and an additional six MDI items with a significance value of one.[273]

The specific content of each symptom complex, importantly, was derived from population health data as collected by the CMI (and later the MDI). Thus, by collecting questionnaire data on thousands of patients, by synthesising and categorising these data according to different variables, and by creating an algorithm for comparing an individual patient's questionnaire responses to population data, Brodman and his colleagues had created a new way of defining and diagnosing disease—one that was statistical and aggregative. The CMI and the MDS therefore did not just standardise medical history taking or discipline clinical observation, they also facilitated new disease concepts and new ways of thinking about patients and patient populations in relation to these concepts. As much as these methods reconfigured the act clinical observation, they also reconfigured the object of that observation: patients and their diseases.

Who decided the contours of the symptom complex? Who chose which symptoms mattered, and how much? To these questions Brodman gave a provocative answer: the computer. 'Our study shows', he asserted in 1961 draft grant proposal, 'that with proper programming, a machine can develop *for itself* a hierarchy of diagnostic significance for symptoms in specific diseases … that proves to be effective in the

[273] List of significant items for the 100 diseases, KBP/3/9.

screening of patients for these diseases'.[274] This was a claim that Brodman had been making since 1959. In his first publication on the machine method published that year, Brodman referred to the 'syndromes [of symptoms] "discovered" independently by the machine…'.[275]

Such language put Brodman in sharp contrast to other leaders in the fields of medical electronics and biomedical computing. In a widely distributed 1960 report on medical electronics, for example, the biomedical computing pioneer Robert Ledley stipulated that 'rigorous diagnostic criteria for the various diseases must be established by medical investigations before computer programming can be meaningful'.[276] 'This attitude', Brodman wrote in his notes, 'is contrary to mine. I believe that the computer should be programmed to establish its own criteria. Our research shows that it is indeed possible to do so and that the machine criteria transcend in usefulness for the machine operation those extablished [*sic*] by the accumulated experience of the medical profession'. By Brodman's rendering, the computer is almost an agentic participant in knowledge-making process, 'discovering' on its own the symptom complexes for each of the 100 diseases. Of course, the computer was not actually determining diagnostic weights on its own: it simply applied a formula (that Brodman and van Woerkom had worked out, to some degree 'by hand analysis') to an existing set of data (that originally existed in paper form). But Brodman's sentiments nicely illustrate the extent to which the IBM 704 allowed researchers to engage with and manipulate large amounts of data in ways that, even if not inconceivable by hand, would be immensely (if not prohibitively) time-consuming and labour-intensive.

[274] Draft of Proposal to Veterans Administration, 28 Feb. 1961, KBP/3/12.
[275] Brodman et al., 'Interpretation of Symptoms', p. 120.
[276] Brodman to Kampelman, 12 Jan. 1960, KBP/3/12.

Brodman and his colleagues were not the only actors interested in how the CMI and the MDS constructed disease categories. Others who encountered this work also speculated about the relationship between 'symptom complex' and 'disease'. Was the symptom complex equivalent to, indicative of, or subsumed by its associated disease? In 1966, a reviewer of one of Brodman's first MDS manuscripts asked for clarification on the nature of the symptom complex: 'What is the definition of disease? Did all the physicians in the study agree that the symptoms complexes detected by the medical data screen were really indicative of disease that should have been detected by the physician? Or did some physicians feel that those symptoms complexes were really detected by the physicians but the physician did not feel they constituted diseased and were therefore not reported?'[277] Little did the reviewer know that he was posing a series of questions that Brodman himself was grappling with. Brodman seemed to believe that the symptom complex, while not equivalent to disease, was suggestive of disease. He frequently pointed out that '[t]here is no one-to-one relation between disease and symptom' and even included as a general warning on all MDS outputs the statement: 'THE MEDICAL DATA SCREEN IDENTIFIES SYMPTOM COMPLEXES; IT DOES NOT MAKE DIAGNOSES. THIS PATIENT MAY HAVE A DISEASE WITHOUT REPORTING ITS SYMPTOMS OR MAY REPORT ITS SYMPTOMS WITHOUT HAVING THE DISEASE'.[278]

Another disclaimer on the MDS output similarly drew attention to the difference between disease and symptom complex: 'THE MEDICAL DATA SCREEN (MDS) HAS MATCHED THIS PATIENT'S SYMPTOMS ON THE MDI QUESTIONNAIRE WITH COMPLEXES OF SYMPTOMS FOUND IN 100

[277] Katharine Boucot to Brodman, 28 Dec. 1966, KBP/4/2.
[278] Entries for MDS Output, n.d., KBP/2/9.

DISEASES. THE MDS IDENTIFIED SYMPTOMS COMPLEXES THAT OCCUR IN THE DISEASES NAMED BELOW. THESE IDENTIFICATIONS BECOME PROVEN DIAGNOSES WHEN CONFIRMED BY EXAMINATION OF THE PATIENT'. Here, Brodman suggests that what distinguishes diseases from symptoms complexes is the physician's proverbial 'touch': a symptom complex only becomes a real disease when the physician says so. Although symptom complexes as operationalised by the MDS were quite different from conventional definitions of disease, for Brodman the crucial difference had less to do with any properties internal to these concepts and more to do with their source—with *who* (or *what*) was doing the diagnosing. A machine diagnosed symptom complexes; a human diagnosed diseases. Brodman's efforts to drive a conceptual wedge between 'disease' and 'symptom complex' therefore represented a way of signaling that, even with the introduction of computers into diagnosis, the real work of medicine remained the province of human physicians.

Brodman clearly went out of his way to distinguish symptom complexes from diseases, but in other writings he often slipped back and forth between the two. While easy to discriminate between the categories conceptually, the distinction proved difficult to maintain in practice. Frequently Brodman referred to the two concepts interchangeably. For example, in a 1967 letter, he contradicted earlier assertions that the MDS method identified symptom complexes as opposed to diseases: 'The MDS identifies diseases unrecognized by physicians…'.[279] This was not a one-off; in a grant application filed a decade earlier, Brodman said of the MDS, 'The method correctly diagnosed those diseases about which the questionnaire elicited diagnostically useful

[279] Brodman to Wander, 27 Sep. 1967, KBP/4/2.

data'.[280] These statements about the role of the MDS in diagnosing disease depart from the careful declarations made elsewhere that the MDS only identifies symptom complexes. These inconsistencies even pervaded published materials like the MDS Manual. In revising the Manual in 1971, Brodman toggled between 'disease' and 'symptom complex' (Figure 7). Such confusion may have prompted Brodman and a collaborator to consider jettisoning the entire concept of symptom complex. As he noted in a 1971 letter to Marvin Miller, 'Ralph and I have been considering many matters, including the question whether we should cease using the term "symptom complex," in referring to a disease identified by MDS, or to use instead a term such as, "presumptive disease," or "possible disease"'.[281]

Whatever the nature of the symptom complex there was little question that the MDS put the symptom front and centre. This privileging of symptoms ran counter to trends in mid-twentieth-century medicine. Due to the growing reach of laboratory medicine, symptoms were losing ground to signs. That is, 'subjective' measures of health—markers of disease that only the patient could discern—were relegated to a secondary role in comparison to 'objective' measures of health—markers of disease that a physician could discern independent of the patient's experience.[282] Consider diabetes. Once made entirely on the basis of symptoms like polyuria (increased urination), polydipsia (increased fluid intake) and polyphagia (increased appetite), the diagnosis of diabetes was increasingly made on the basis of haemoglobin A1C tests,

[280] Keeve Brodman, 'A screening method for the 100 most common diseases', Application for Resarch Grant (RG-6302), National Institutes of Health, Division of Research Grants, 31 Oct. 1958, KBP/3/12.
[281] Brodman to Miller, 18 Apr. 1971, KBP/5/1.
[282] Physicians at the time associated symptoms with the subjective and signs with the objective. Walter Modell, 'See The Full Treatment—A Modern View of the Relief of Symptoms', *NEJM* 255 (1956): 1079-1084.

urine tests, blood glucose tests, and glucose tolerance tests.[283] Indeed, the move away from symptoms permitted the emergence of an entirely new disease category: asymptomatic pre-diabetes. Whereas the patient's symptomology represented the surface expression of disease, many hailed these kinds of laboratory tests as windows into the etiological mechanism of disease. The growing reliance on the pathognomonic sign occasioned a shift in the logic behind the doctor-patient relationship—from 'trust' to 'trust, but verify'.

In defending his method's reliance on symptoms over signs, Brodman invoked the accumulated wisdom of centuries of medical practice. The diagnostic significance of symptoms seemed self-evident to Brodman; even in age of laboratory medicine physicians attend to patients' symptoms, just as they had for centuries. As Brodman put it curtly in one grant application: 'Physicians recognize the diagnostic value of symptoms'.[284] The defence of symptoms did not come with the derision of signs; Brodman saw the two as mutually supportive. In fact, early on in the creation of the MDS, Brodman expressed optimism that the method would expand over time to 'one day include data from physical and laboratory examinations' alongside the original data derived from symptoms.[285] Once the MDS method actually began to take shape, however, he moved away from this early view and came to the opinion that their computer method ought to remain separate from—but also complementary to—other laboratory examinations. The MDS team incorporated this complementarity into the MDS. The MDS output, in addition to urging the physician generally to confirm the provisional diagnosis with 'appropriate data obtained in … laboratory tests', also

---

[283] Jeremy A. Greene, *Prescribing by Numbers: Drugs and the Definition of Disease* (Baltimore: Johns Hopkins University Press, 2007), pp. 230-231.
[284] Keeve Brodman, Application for Research Grant (RG-4743), National Institutes of Health, Division of Research Grants, 20 Jan. 1960, KBP/1/7.
[285] Brodman to van Woerkom, 16 Feb. 1960, KBP/6/3.

included a section—the 'Physician's Order Sheet'—that listed specific laboratory examinations that might be helpful in determining the presence and severity of the diseases suggested by the MDS.[286]

By conventional accounts, defining disease in relation to symptoms sits in opposition with defining disease in relation to numerical thresholds. Glucose levels, breathing rates, blood pressures, pulses—these were easy to quantify. But a patient's description of his own condition? Not so much. The MDS method, however, bound these two ways of defining disease—symptomologically and numerically—together. The MDS method turned a given patient's questionnaire responses (his symptomology) into different numerical scores, one for each of the diseases considered by the method. To calculate these scores, the MDS used significance values derived from the patterns in the symptomology of the larger populace. To convert these diagnostic scores into an estimation of diagnostic likelihood, the MDS compared the score of a particular patient against the mean score of all patients who were confirmed as having the disease in question (Figure 8). The MDS method thus made the larger population—or at least a supposedly representative sample of the population—the ultimate reference point for health and disease. By attaching a numerical score to patterns of symptoms, one might interpret that MDS as what Steven Shapin has recently called an 'objectivity engine': it transformed subjective experiences (symptoms) into a form that could be readily measured, quantified, manipulated, compared, and communicated.[287] In short, the MDS defined pathology both symptomologically and numerically.

---

[286] Entries for MDS Output, n.d., KBP/2/9; 'MDS Output, for Manual, page 33, after the 7th life down', n.d., KBP/2/10.
[287] Steven Shapin, 'A Taste of Science: Making the Subjective Objective in the California Wine World', *Social Studies of Science* 46, no. 3 (2016): 436-460, p. 451.

The case of the CMI and MDS suggests that the process of disease (re)definition is not just a social or cultural process. It is a material process as well. The rise of 'the statistical patient' was facilitated by the use of computers and the digitisation of questionnaire data. But even more importantly, the MDS method's use of statistics was enabled by the standardised nature of paper questionnaires. Underlying these statistical characterisations of disease, then, were *paper* technologies—medical questionnaires that standardised medical histories and made them amenable to mathematisation, compilation, and ultimately computerisation. Brodman pointed out the tight connection between the standardised questionnaire and the statistical definition of pathology in a 1960 grant application: 'Before the CMI made possible the collection of a large body of medical data from many patients, statistical diagnostic analysis was impracticable'.[288] Understanding how the MDS promoted certain definitions of disease requires an understanding of the material culture of medicine.

* * *

The aspirations that Brodman had for the MDS remained just that—aspirations. In December 1972, Miller wrote to Brodman availing him of Roche's most recent market survey results: 'We find that the most recent survey findings challenge the successful marketability of MDS. Further, it appears that the market may be smaller than was originally anticipated'. Because Roche's primary interest in the MDS was financial, these new insights into the limited profitability of the programme sounded the death knell: 'Regrettably, as a matter of reasonable business judgment, we find it necessary to terminate both agreements between Roche and yourself'. The news came as 'a shocking surprise' to Brodman. It apparently also surprised Roche personnel,

[288] Keeve Brodman, Application for Research Grant (RG-4743), National Institutes of Health, Division of Research Grants, 20 Jan. 1960, KBP/1/7.

who were actively working on edits to the MDS as late as one week before Miller's notice. The revelation came as Brodman's career was already in its eclipse. Brodman had been retired from Cornell since 1971, and since the late 1960s he found his multiple sclerosis to be almost 'incapacitating'. Roche's withdrawal from work on the MDS therefore wrested from Brodman his last remaining academic engagement.

The MDS is of far greater significance than its ultimate fate might suggest. Above all, the history of the MDS suggests the need for a new genealogy of computerised diagnosis, one that makes cybernetic thinking and paper data practices more central. Relatively little has been said about the role of cybernetic thinking in the history of biomedical computing; while most of the actors covered in the historiography were surely aware of cybernetics, these pioneers in biomedical computing were not grappling with cybernetic themes and principles in their actual work. Brodman was.[289] It was by way of cybernetics that Brodman, a psychiatrist, came to think about questions of mathematisation, mechanisation, and computerisation. His cybernetically inflected efforts to model clinical cognition bridged his earlier paper-based work and his later computer-based work.

Brodman's case also highlights how the history of computerised diagnosis has its deepest roots not in electronic digital computers but in earlier paper-based technologies—and the epistemologies and material practices associated with these paper tools. The mathematics behind the MDS was created with respect to data that had originally been collected by a paper instrument for the purpose of evaluating that paper instrument. The process of devising and testing these statistical formulas, moreover,

[289] Heinz von Foerster, head of the Biological Computer Laboratory (BCL), stands as another example. Albert Müller and Karl Müller (eds.), *An Unfinished Revolution? Heinz von Foerster and the Biology Computer Laboratory (BCL) 1958-1976* (Vienna: Edition Echoraum, 2007).

involved a hybrid of automated and manual modes of tabulation. Once finally developed and formally proposed, the computerised MDS method continued to rely upon paper questionnaires and paper printouts and the physical circulation of these documents among different actors. In short, the CMI and the MDS were intimately intertwined: there is no telling the history of the computer-based MDS without also telling the history of the paper-based CMI. The evolution from CMI to MDS was less a sharp transformation than a gradual, multi-year transition from one technology to another. The contingent decisions that Brodman and his colleagues made when creating the CMI—the questions they chose to ask and the reams of data they would collect—ultimately constrained the form and the function of the MDS many years later.

This section, very importantly, does not reject the idea that computers have brought changes, sometimes profound, to medicine and medical practice. Indeed, while this section has made the case for the continuity between paper-based and computerised practices, it has also pointed to changes that accompanied the introduction of computing technologies to Brodman's work. Namely, computing technologies greatly reduced the 'friction' associated with questionnaire-derived data, allowing researchers to process, analyse, understand, and manipulate these data with much greater ease. Recognising such changes, however, remains consistent with my broader point that these data—and ideas about how they might be used—evolved directly from earlier paper-based practices. This section, then, asks that as historians, science and technology scholars, and others work to identify and understand technological change, they not lose sight of the continuities—of the ways in which the medical field's earliest engagements with computers may have been shaped, constrained, and enabled by earlier practices, cultures, and epistemologies. Acknowledging these continuities opens

up new possibilities for engagement among historians working in different time periods and across different material cultural traditions.

PART II

Recoding Disease

- CHAPTER III -

'The Standardized Language of the Machine'

From the perspective of most residents, 10 December 1957 was a typical day in Camden, New Jersey (Figure 9). But from the perspective of an interdisciplinary group of researchers based out of the Rockefeller Institute, the day was anything but. Beginning at 11 a.m., this research team—comprising physicians, statisticians, and engineers and headed by the television pioneer Vladimir Zworykin—gathered at the Radio Corporation of America's (RCA) laboratories in Camden to witness a demonstration of what may have been the first instance of a computer diagnosing disease in U.S. history: they watched as the RCA's Bizmac computer rendered, from inputs of patient data, diagnoses of haematologic diseases (Figure 10).[290] Recognising the significance of this moment, the researchers organised around the demonstration a series of speeches, discussions, and panels, and they invited a select group of leading physicians and engineers with an interest in the emerging field of medical electronics to participate.

Others registered the significance of this effort too. In 1965, the *New York Times* ran the headline 'Computer Now Can Diagnose Ills', announcing that Zworykin's group had 'invented an automatic system for diagnosing disease' and had recently filed for a related patent.[291] The demonstration in Camden was less of an end than a beginning: the haematology programme carried on for nearly three decades, evolving in its aims, methodologies, and group membership. This section reconstructs

[290] 'A Bizmac Program for Medical Data Processing', December 10, 1957, REP/14/10.
[291] Stacey V. Jones, 'Computer Can Now Diagnose Ills', *New York Times*, 6 Nov. 1965, p. 35.

this effort—one that, despite its significance, makes only a blip in the existing historical record.

The haematology programme offers a window into the role of the computer as an agent of disease definition. This section argues that physicians and engineers, in creating computer systems that screen for and diagnose disease, simultaneously shaped the meanings and definitions of diseases themselves. More specifically, it charts how the introduction of computers to medical diagnosis both resulted from and promoted disease concepts that were ultimately numerical and probabilistic. This process of disease redefinition, I submit, was a material one—the result of having converted medical knowledge into different material forms. As Zworykin described it, this was the process of 'translat[ing]' medical knowledge and patient information 'into the standardized language of the machine'.[292]

Existing scholarship on the definitions of disease foreground the social, intellectual, and cultural factors that give rise to different disease concepts. From venereal disease in the nineteenth century to diabetes in the twentieth, much has been said about how historically-contingent moral attitudes, intellectual tools, social arrangements, clinical behaviours, and marketing practices have shaped our understandings and definitions of disease. But as this section seeks to make clear, the process of disease (re)definition is not just a social or intellectual process. It is a material process as well, shaped profoundly by the material media with which physicians work. Charles Rosenberg has noted that '[c]oncepts of disease … always

[292] Vladimir Zworykin, 'New Frontiers in Medical Electronics: Electronic Aids for Medical Diagnosis', draft paper, 27 Sep 1963, box 78, folder 59, Vladimir K. Zworykin Papers, 1908-1981, David Sarnoff Research Center records (Accession 2464.09), Manuscripts and Archives Department, Hagley Museum and Library, Wilmington, Delaware (hereafter cited in the form VZP/78/59).

exist in both social and intellectual space'.[293] As this section seeks to make clear, disease concepts exist in material space as well. Definitions of disease don't just have a history. They have a body as well—they are incarnate in stuff.

THE BAYESIAN BANDWAGON

At the same time Zworykin was launching his work on the haematology project, another team—the duo of Robert Ledley and Lee Lusted—were just beginning a fruitful intellectual collaboration. This collaboration ultimately elaborated a strategy of automating diagnosis that attracted an exceptional level of attention. The strategy revolved around Bayes' theorem. In 1959, this pair of researchers published in the journal *Science* their ground-breaking article 'Reasoning Foundations of Medical Diagnosis'.[294] Although the article gestured to the use of digital electronic computers in medical diagnosis, its immediate aim was more technical: to outline different methods of rendering diagnostic reasoning in mathematical terms. The article addressed many intellectually delicate questions. What is the optimal way of representing, with symbolic logic, the relationship between combinations of symptoms (symptom complexes) and combinations of diseases (disease complexes)? How can one mathematically analyse treatment decisions given that a patient's choice of treatment 'not only depends on the established diagnosis but also on therapeutic, moral, ethical, social, and economic considerations'?[295] The article's most enduring intervention, however, was its application of the well-known Bayes' theorem to medical diagnosis.

---

[293] Rosenberg, 'Framing Disease', p. xx. See also Charles Rosenberg, 'The Tyranny of Diagnosis: Specific Entities and Individual Experience', *Milbank Quarterly* 80, no. 2 (2002): 237-260.
[294] Robert Ledley and Lee Lusted, 'Reasoning Foundations of Medical Diagnosis', *Science* 130, no. 3366 (Jul. 1959): 9-21.
[295] Ibid., p. 15.

Ledley and Lusted proposed that a physician, provided with a large enough set of patient data, could calculate the probability of an event (i.e., the diagnosis) based on prior knowledge about conditions that might be related to that event (i.e., symptoms, demographic factors, etc.).

Ledley and Lusted's interest in Bayes' theorem was very much a reflection of their times. The theorem's origin dates back to the mathematical work of the eighteenth-century English minister Thomas Bayes.[296] After enjoying an early period of critical engagement, Bayes' theorem lapsed in and out of relative obscurity over the ensuing centuries. By the early twentieth century, many statisticians believed that Bayesian methods fizzled out essentially from a lack of oxygen: few scholars gave the theorem the thought and attention necessary for its survival and acceptance. But World War II, and the Cold War that followed, changed the landscape of statistics and gave the theorem new life. Alan Turing used Bayes' theorem while working as a code breaker at Bletchley Park; the U.S. Navy relied on Bayesian statistics to help locate missing vessels; the RAND Corporation found the theorem useful in their efforts to estimate the probability of a potential nuclear event.

Within the academy, a group of scholars based predominantly at the University of Chicago and Harvard promoted Bayesian methods as part of what some scholars have called the 'neo-Bayesian revival' of the 1950s and 1960s.[297] The leaders of this revival, Jimmie Savage, Dennis Lindley, and I. J. ('Jack') Good (who worked as a cryptologist with Turing), argued for an approach to statistical inference in which

---

[296] Sharon McGrayne, *The Theory that Would Not Die* (New Haven: Yale University Press, 2011). On the contributions of Pierre-Simon Laplace to Bayesian probability, see Lorraine Daston, *Classical Probability in the Enlightenment* (Princeton: Princeton University Press, 1988).

[297] Stephan E. Fienberg, 'When Did Bayesian Inference Become "Bayesian"?', *Bayesian Analysis* 1, no. 1 (2006): 1-40.

decision makers use Bayes' theorem to update an initial state of knowledge as more information and evidence becomes available. Soon, this interest in Bayesian statistics within the military and the academy spilled over into a number of applied domains. Medicine was no exception.

Ledley and Lusted's professional trajectories in many ways mirrored the life course of Bayes' theorem. As Bayesian methods migrated from military to medical contexts, so Ledley and Lusted arrived at the topic of computers in medical reasoning by way of their military backgrounds. As a student at Columbia University in the 1940s, Ledley was strongly drawn to physics, in part because the school's physics department was then populated by luminaries like Hans Bethe, Enrico Fermi, and Isidor Rabi. Family expectations, however, pulled him in a different professional direction. Ledley's father worried that physics offered few professional opportunities and encouraged him to go into a more practical field, like dentistry.[298] Ultimately Ledley did both, undertaking graduate work in physics at Columbia while completing his dentistry degree at New York University. In 1950, as the U.S. was ramping up its involvement in the Korean War, Ledley joined the U.S. Army Dental Corps, allowing him to conduct dental research at Walter Reed General Hospital in lieu of serving as an infantryman abroad.

Ledley's military research in dental prosthetics resulted in a series of collaborations with—and later a job at—the National Bureau of Standards (NBS), where he would encounter the Standards Eastern Automatic Computer (SEAC), a first-

[298] Robert S. Ledley, 'Medical Informatics: A Personal View of Sowing the Seeds', in B. Blum and K. Duncan (eds.), *A History of Medical Informatics: Proceedings of the Association for Computing Machinery Conference on History of Medical Informatics* (New York: ACM Press, 1990). Dean Sittig, Joan Ash, and Robert Ledley, 'The Story Behind the Development of the First Whole-body Computerized Tomography Scanner as Told by Robert S. Ledley', *Journal of the American Medical Informatics Association* 13, no. 5 (2006): 465-69.

generation digital electronic computer. It was this introduction to computers that reoriented his interests to questions of biomedical computing. '[T]he SEAC', he wrote, 'was my panacea'—the computer gave him hope that mathematical and physical formulas could be useful in solving, understanding, and representing complex biomedical phenomena (Figure 11).[299] In 1954, following a wave of deep funding cuts, Ledley left NBS and took up a research position in the Operations Research Office at Johns Hopkins University. There, Ledley worked on applications of formalised mathematical logic to complex problems of military strategy. Ledley's foray into the computerisation of medical reasoning represented the coming together of these many life strands: his background in medicine, his interest in computing, and his expertise in operations research (OR) techniques.

When Ledley was in Baltimore investigating OR applications to biomedicine, Lusted was in Bethesda embarking on research in biomedical computing at the NIH. Lusted, though trained as a medical doctor, was well versed in electrical engineering. Between graduating from college in 1943 and starting medical school in 1946, he worked briefly for Harvard's Radio Research Laboratory, installing radar in American fighter aircrafts and studying Germany's radar technology. Nevertheless, his arrival at the NIH and his computational research there represented a deviation from his earlier professional trajectory. In 1955, as Lusted was laying the groundwork for a career in academic radiology in the San Francisco-Bay Area, the U.S. Government presented Lusted with an ultimatum similar to the one Ledley faced five years before: either serve active duty for two years in the PHS (where Lusted had been on reserve commission), or risk being drafted into the Army. In opting for the former, Lusted eventually ended up at the NIH.

[299] Ledley, 'Medical Informatics', p. 90.

For someone interested in electronics, the mid-1950s was a stimulating time to be in Bethesda. During this period, the NIH initiated, in partnership with outside companies, a range of projects geared toward developing new medical electronic devices. Lusted accepted an invitation from the Airborne Instruments Laboratory (AIL) to consult on the development of one of these projects, the Cytoanalyzer. Produced under the direction of the AIL physicist Walter Tolles, the Cytoanalyzer was an optical electronic machine designed to aid in the detection of uterine cancer by automated computer analysis of vaginal ('Pap') smears.[300] Lusted, though enthusiastic about the cytoanalyzer and its medical potential, was left a little longing; what the field still lacked, Lusted felt, was a broader understanding of how computational techniques might assist the more fundamental processes of medical diagnosis and reasoning. It was while in this frame of mind that Lusted first encountered Ledley's research through a presentation that Ledley gave at the 1956 annual meeting of the Operations Research Society of America.[301] With the aid of a mutual friend, the prominent OR researcher William Horvath, Lusted made contact with Ledley, touching off a series of conversations about computers and their potential uses in biomedicine and medical diagnosis. Out of these conversations their paper, 'Reasoning Foundations of Medical Diagnosis', was born.

Much to their surprise, this 1959 paper (despite its quiet and circumstantial origins) attracted an enormous audience, one that comprised groups as various as computer specialists, medical practitioners, mathematicians, logicians, and lay people.

---

[300] 'Cytoanalyzer', *Public Health Reports* 72, no. 11 (1957): 1038. On the cytoanalyzer from an STS perspective, see Peter Keating and Alberto Cambrosio, *Biomedical Platforms: Realigning the Normal and the Pathological in Late-Twentieth-Century Medicine* (Cambridge and London: MIT Press, 2003), pp. 63-66.
[301] Rudolf Seising, 'Robert S. Ledley, 1926–2012', *Artificial Intelligence in Medicine* 57, no. 1 (2013): 1-7.

Their article quickly circulated among biomedical researchers, physicians, and computer scientists, through both personal correspondence and professional journals.[302] As Ledley later recalled, the paper 'stimulated a tremendous amount of work in these areas [medical informatics] and has been called the seminal paper starting a lot of this work. It's very hard to believe, looking back, but there's no way I could have visualized at that time how influential that paper would turn out to be'.[303] A number of media outlets further amplified the article's reach. The *New York Times*, *New York World Telegram*, and *New York Post* all used Ledley and Lusted's work as a springboard into speculative discussions about the future of computers in medicine.[304]

A number of research centres translated the theoretical sketch offered by Ledley and Lusted into functioning diagnostic systems. The most prominent of these was based out of the University of Utah and the Latter-day Saints (LDS) Hospital in Salt Lake City.[305] There, Homer Warner, together with his colleagues Allan Pryor and Reed Gardner, began building what would arguably become the most successful health information infrastructure in the country.[306] The project started small. Inspired by Ledley and Lusted's mathematical formalisations, Warner and his colleagues grew interested in the possibility of applying ideas about Bayes' theorem and conditional

---

[302] Through the 1960s, and even into the 1970s, pioneers in biomedical computing discussed the article frequently and circulated it widely. See, e.g., letter from Joshua Lederberg to Luca Cavalli-Sforza, 25 Oct. 1960, The Joshua Lederberg Papers, National Library of Medicine, National Institutes of Health, Bethesda, Maryland, box 11, folder 81 (hereafter cited in the form JLP/11/81).
[303] Ledley, quoted in 'Planting the Seeds: Panel Transcript', in *A History of Medical Informatics,* p. 61.
[304] November, *Biomedical Computing*, pp. 62-64.
[305] For more on this history, see Joseph A. November, 'Early Biomedical Computing and the Roots of Evidence-Based Medicine', *IEEE Annals of the History of Computing* 33, no. 2 (2011): 9-23.
[306] These efforts merged with the Intermountain Healthcare system, which in November 2009 was cited by President Obama as a model for health care reform. David Leonhardt, 'Dr. James Will Make It Better', *New York Times Sunday Magazine,* 8 Nov. 2009, p. MM31.

probabilities to actual clinical data.[307] They quickly settled on congenital heart disease as their test domain. Their reasons for selecting this domain were twofold: data quality and data quantity. Quality-wise, the accuracy of diagnoses made on the basis of clinical findings could be verified against the gold standard of cardiac catheterisation and findings at surgery. Quantity-wise, the team believed that 'relatively objective clinical findings' could be easily obtained. The team thus assembled a database: a matrix that included incidence information on about 57 symptoms and 35 diseases.

Warner translated this mathematical approach for implementation on a digital electronic computer.[308] The computer's performance impressed Warner—and many others. The group found that, in a series of 74 patients with congenital heart disease, the diagnoses made by computer were essentially identical to those made by the physicians who had seen the same patients. The results grabbed national attention. 'Computer Is Found Useful in Heart Disease', the *New York Times* reported. 'A computer has been programmed so that it can diagnose congenital heart defects as accurately as can an experienced heart specialist. In fact, it was reported, physicians working along the computer improved their own diagnostic skill while the computer improved its own ability to pinpoint heart defects'.[309] Excitement around the findings fuelled familiar speculation about the technological deskilling of the profession, with '[c]omputers … tak[ing] the place of highly skilled specialists, and more physicians … becoming well-

[307] Homer R. Warner, Alan F. Toronto, L. George Veasey, and Robert Stephenson, 'A Mathematical Approach to Medical Diagnosis: Application to Congenital Heart Disease', *Journal of the American Medical Association* 177, no. 3 (1961): 177-183.
[308] Alan F. Toronto, L. George Veasy, and Homer R. Warner, 'Evaluation of a Computer Program for Diagnosis of Congenital Heart Disease', *Progress in Cardiovascular Disease* 5, no. 4 (1963): 362-377.
[309] Robert H. Plumb, 'Computer Is Found Useful in Heart Diagnoses: Machine Taught to Analyze Symptoms Accurately, 3 Physicians Report', *New York Times*, 28 May 1963, p. 19.

trained data-collectors'.[310] The findings also put wind in the sails of Warner's larger research enterprise at Utah. Before long, Warner's narrow focus on cardiologic diagnosis broadened. The project that started in the Cardiovascular Laboratory soon spread outward, evolving into a larger health information system called Health Evaluation through Logical Processing (HELP). By the late 1960s, HELP had been adopted widely across the LDS Hospital.[311]

Yet, the growing excitement around Bayes' theorem was tempered by growing realisations about its limitations. The haematology team were highly tuned into these limitations. Carl Berkley described the team's abnegation of Bayes' theorem in a 1966 letter: 'In the evaluation of various problems of diagnosis we have been led to a detailed examination of diagnostic processes and to a realization of the inherent pitfalls in total reliance of Bayesian methods'.[312] Max Woodbury, a statistician at New York University who would come to lead the haematology project for a period of time following Zworykin's true retirement in 1961, unpacked some of the 'inherent pitfalls' to which Berkley was referring.[313]

In a 1963 talk delivered at the Fifth International Conference on Medical Electronics, Woodbury discussed the mathematical and logical basis behind what he saw as the 'Inapplicabilities of Bayes' Theorem to Diagnosis'.[314] Woodbury emphasised the assumptions that were required by Bayes' theorem yet violated by

---

[310] Ibid.
[311] On the history of HELP and biomedical computing at Utah and LDS Hospital, see, Gilad J. Kuperman, Reed M. Gardner, and T. Allan Pryor, *HELP: A Dynamic Hospital Information System* (New York: Springer-Verlag, 1991), pp. 3-13.
[312] Carl Berkley to Louis Gerstman, 9 Mar. 1966, REP/14/2.
[313] For other discussions about the disadvantages of Bayesian approaches, see, e.g., Eli Robins, 'Categories Versus Dimensions in Psychiatric Classification', *Psychiatric Annals* 6, no. 8 (1976): 39, 42-43, 46-47, 51, 55.
[314] Max A. Woodbury, 'Inapplicabilities of Bayes' Theorem to Diagnosis', in C. C. Thomas (ed.), *Proceedings of the Fifth International Conference on Medical Electronics*, Liege, Belgium (Springfield, IL, 1963), 860-868.

medicine. ‘The primary difficulty in the application of the theorem’, Woodbury explained, ‘is that it has been assumed that diseases are mutually exclusive and that diseases outside of the list assumed do not produce symptoms considered in conjunction with the diagnosing of diseases in the group’.[315] In short, Bayes theorem assumed that each symptom is independent, that each diagnosis has an equal probability, and that each diagnosis is independent of the others. Such assumptions ran afoul of medical reality—almost to a laughable degree. Of course patients can have multiple diseases at once; of course certain clusters of symptoms run together; of course the diseases considered by a diagnostic system are not exhaustive of all diseases that patients could possibly have. Indeed, such Bayesian violations may have been the rule, not the exception, for many patient populations.

Besides faulty assumptions, the haematology team also worried about Bayes’ theorem’s data requirements. The theorem depended upon the existence of accurate conditional probabilities. Yet, apart from places like Utah and diseases like congenital heart disease, few areas of medicine were sufficiently data rich to derive these probabilities from population data. For a computerised diagnostic system to have any generalisability, the haematology team believed, a different approach was needed.

RECODING DISEASE

The demonstration in Camden was the culmination of some five years of work. Towards the beginning of September in 1954, a lavish event was held to honour Vladimir Zworykin in the wake of his retirement from the RCA. Then 64 years old, Zworykin was a towering figure not just in the world of engineering but in American intellectual and cultural life more generally, having invented the iconoscope and the

[315] Ibid.

cathode ray tube that made modern electronic television possible (Figure 12).[316] Attendees who knew Zworykin well predicted that retirement, in his case, wasn't so much an end as it was a new beginning. As RCA president David Sarnoff declared at the event, 'I see no relationship whatever between retirement and Vladimir Zworykin. ….[H]e will continue to do whatever work interests him most'.[317] Sarnoff was right, and what began to interest Zworykin most was the nascent field of medical electronics.

Around the time of his retirement, Zworykin grew interested in the work of two physicians, Martin Lipkin and James Hardy. Since 1952, Lipkin and Hardy had been studying medical decision making and developing a mechanical method of using marginal punched cards to aid in the differential diagnosis of hematologic disease. The device they proposed was ingenious in its simplicity.[318] For each disease considered, the researchers prepared a single 'master card'. Punched onto the periphery of every master card were holes, each of which corresponded to a different piece of clinical information. For a given master card, triangular wedges were then punched into a unique set of marginal holes, marking the clinical data that were characteristic of that particular disease (Figure 13). Once prepared, the master cards for each disease were collected and arranged into a deck. Diagnosing a given hospital case involved inserting

---

[316] On Zworykin in the history of television, see Albert Abramson, *Zworykin, Pioneer of Television* (Urbana: University of Illinois Press, 1995); Albert Abramson, *The History of Television, 1880-1941* (Jefferson, NC: McFarland, 1987); Alexander Magoun, *Television: The Life Story of a Technology* (Westport, CT: Greenwood Press, 2007).

[317] Address of Brig Gen. David Sarnoff, Testimonial Dinner and Seminar in Honor of Vladimir K. Zworykin, VZP/87/5.

[318] Martin Lipkin and James Hardy, 'Differential Diagnosis of Hematologic Diseases Aided by Mechanical Correlation of Data', *Science* 125, no. 3247 (1957): 551-552; Martin Lipkin and James Hardy, 'Mechanical Correlation of Data in Differential Diagnosis of Hematological Diseases', *JAMA* 166, no. 2 (1958): 113-125. The work of Lipkin and Hardy built upon the earlier work of Firmin A. Nash on the 'Logoscope', a diagnostic slide rule. See F. A. Nash, 'Diagnostic Reasoning and the Logoscope', *Lancet* 276, no. 7166 (1960):1442-1446.

a rod through the punched hole corresponding to a clinical finding present in that case. After the rod was inserted and raised, only disease cards without triangular wedges at that specific punched hole (i.e., diseases for which that clinical finding was *not* characteristic) were raised and removed from the deck. By repeating this process for additional clinical findings in the hospital case, a physician could continually reduce the number of disease cards that matched the patient's clinical picture. Eventually, the physician would arrive at just a few cards. Among these cards, Lipkin and Hardy hoped, would be the patient's most probable diagnosis.

Zworykin saw in Lipkin and Hardy's work an opportunity. He reached out to them and pointed out that their mechanical strategy of matching patient data with disease descriptions was amenable not just to paper methods but to electronic computing procedures as well. The same year that he retired from RCA, Zworykin had secured a new perch from which to explore this possibility. As the director of the Medical Electronics Center at the Rockefeller Institute, he oversaw various research projects in the area of medical electronics. Working at the Rockefeller Institute proved highly conducive to the pursuit of exploratory, novel, and ambitious research projects of the kind that interested Zworykin. The Rockefeller Institute stood apart from other academic institutions in its decentralised arrangement. Organised around individual researchers rather than departments, the Institute provided researchers like Zworykin with more latitude to pursue research in any area, irrespective of how this research might breach conventional disciplinary and departmental boundaries.[319] For Zworykin's part, this freedom was particularly important; Zworykin stressed that the field of medical electronics was nothing if not interdisciplinary. 'In the new

[319] *The Rockefeller Institute Bulletin, 1960-1961*, volume 5, no. 2, pp. 36-37, Rockefeller University records, Catalogs and Directories, Rockefeller Archive Center, Sleepy Hollow, New York, box 1, FA207 (hereafter cited in the form RAC/FA207/1).

interdisciplinary field of medical electronics', he explained in a 1961 speech, 'we hope to break down these walls and restore the unity of science in at least one sector of vital importance for human welfare'.[320]

True to this ideal, Zworykin forged an interdisciplinary team of researchers to explore the question of how to computerise medical diagnosis. The team brought together engineers and computing experts (Zworykin, Carl Berkley, Russell Ebald, Robert Lane, and Marvin Sendrow); physicians (Lipkin, a gastroenterologist, along with B. J. Davis, M. A. Atamer, and Ralph Engle, Jr., all haematologists); and a statistician (Max Woodbury). This group convened at the Rockefeller Institute for monthly seminars, during which they discussed virtually all aspects of computer aids to diagnosis. At some meetings their discussions were technical and mathematical. But on other days, they ventured into more philosophical territory, discussing how—and how much—physicians can know.

Wandering into uncharted territory, the research group recognised that they would need to set certain parameters and select a narrow subfield of medicine with which to work and test their ideas. For a number of reasons, haematology quickly emerged as the team's field of choice. One reason was path dependence: haematology was the field that had been selected by Lipkin and Hardy in their earlier work using marginal punched cards. More significantly, though, the researchers believed haematology to be a field of medicine that was uniquely objective, well understood, and quantitative. Whereas the psychiatrist worked with messy concepts like human behaviour and emotion, the haematologist worked with clear and quantitative measures: theirs was a world of blood cell counts, blood sugar levels, sedimentation

[320] Vladimir Zworykin, 'Medical Electronics', draft of presentation at the University of Bologna, October 1961, VZP/78/27

rates, clotting times, and numerical fragility measurements. As the group explained in 1957, 'Hematology has been chosen for the initial study…because many of the findings in hematology are from clinical laboratory data which are relatively independent of subjective findings of either the patient or doctor'.[321]

As the team ventured into the weeds of clinical haematology, Lipkin realised that it would be prudent to recruit haematologists to join the group. In 1955, he reached out to Ralph Engle, a haematologist at Cornell, inviting him 'to join a group of computer specialists and physicians meeting at Rockefeller Institute to discuss and develop ways in which electronic computers might be used to help physicians'.[322] As one of the team's resident haematologists, Engle later quipped that views about the 'objective' nature of haematology reflected less the reality of haematology and more Martin Lipkin's status as a gastroenterologist: 'The other fellow's specialty always seems more straightforward and less complex!'[323] Without the separation of time, Engle's emotions on the matter seemed a little more raw, less playful; in 1960, he registered being 'upset to learn that others thought hematology was so clear-cut and so definite, without any diagnostic problems, that everything could be handled by computer'.[324]

With many boasting about the promises of a diagnostic machine, the clinically-oriented Engle often found himself as the tempering voice. As early as 1960, he was pushing back against the idea that the programme could be viewed as an autonomous

---

[321] 'A Bizmac Program for Medical Data Processing', 10 Dec. 1957, REP/14/10.
[322] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making: A Thirty Year Experience', *Perspectives in Biology and Medicine* 35, no. 2 (1992): 207-219, p. 208
[323] Ibid., p. 209.
[324] Discussion of Martin Lipkin, 'Correlation of Data with a Digital Computer in the Differential Diagnosis of Hematological Diseases', *IRE Transactions on Medical Electronics* ME-7, no. 4 (1960): 243-246, p. 246.

diagnostic entity: ‘There has been some criticism of the program—in particular, that it is developing a diagnosis machine. I think we have to do everything possible to discourage this concept. It really represents a giant textbook of medicine, which stores not only a clear picture of a particular disease, but also such things as bibliographic data and all types of medical information which the physician can then use to improve his practice and to help him with diagnosis; but this is not a diagnosis machine’.[325]

The team’s first computer programme functioned analogously to the mechanical, punched card method that had been outlined by Lipkin and Hardy: a given patient’s findings were compared to ‘characteristic’ descriptions of each haematologic disease. These definitions of each disease were arrived at through a fairly straightforward process. The group’s haematologists assigned weights to different clinical items (i.e., symptoms, physical findings, and laboratory findings) depending on how—and how strongly—the item was associated with the disease. A weight of 0 indicated that the item was typically present in the disease; 1 indicated that the item was necessary but not necessarily sufficient; and 2 indicated that the item was either irrelevant or not typically present in the disease. Mechanistic, organic, or pathophysiological conceptions of disease held little purchase in the computer’s diagnostic process. The definition of each disease, as far as the computer programme was concerned, was a collection of numerical representations of the association between the disease and given clinical variables.

In her work on the history of automated mathematical proof, the historian Stephanie Dick demonstrates how epistemologically interesting things can happen in

[325] Ibid., p. 246.

the process of *implementation*.[326] The material affordances of computers and machines, Dick argues, can affect how knowledge is organised and structured. Mathematicians, engineers, and computer scientists—in the process of moving mathematical proofs from analogue media dominated by pencil and paper and to new electronic and digital media—were forced to alter their approaches and epistemologies to suit new materialities. For example, Dick shows how Allen Newell, John Clifford Shaw, and Herbert Simon, in their work automating theorem proving on the JOHNNIAC computer, introduced new ways of formulating logical propositions (i.e., linked lists) as a way of navigating the JOHNNIAC's material requirements (i.e., limited memory availability). These linked lists rendered logical propositions in a very different form than their pen-and-paper instantiations (a 'concatenated series of symbols').[327]

What Dick observes in the context of automating mathematical proofs also applies, at least in some cases, in the context of automating diagnosis. The haematology research team was forced to reckon with the material constraints of the computer from the outset. Among these constraints was the computer's processing speed. Even with their first program, which considered only twenty-one haematologic diseases, the researchers found that the computer would take too much time if it went about individually checking the degree of correspondence between the patient data and each of the twenty-one disease descriptions. 'One of the problems in computer diagnosis', the team expounded in a 1962 report, 'is to reduce the number of diseases under consideration by the diagnostician and the computer at any one time'.[328] The group found a taxonomic solution to this material problem: they grouped the twenty-one

[326] Dick, 'AfterMath'; Dick, 'Of Models and Machines'; Stephanie Dick, 'After Math: (Re)configuring Minds, Proof, and Computing in the Postwar United States', Doctoral Dissertation (Cambridge, MA: Harvard University, 2015).
[327] Dick, 'Of Models and Machines', p. 631.
[328] 'Detailed Report and Future Plans: January 1, 1962-December 31, 1962', REP/14/2.

diseases into nine classes. By clustering diseases into larger disease classes, the computer could perform a preliminary matching between the patient's findings and the smaller number of disease classes. Only diseases within the subset of classes found to meet some level of correspondence would be investigated further. These groupings, importantly, were brand new: they did not, the team noted in a 1968 grant application, 'conform completely to any recognized medical classification'.[329] The contours of disease categories, in other words, were redrawn.

With descriptions of both disease classes and individual diseases stored in the computer, the last piece of information that needed to be fed into the computer was the clinical data from the patient under consideration. To facilitate this process, the researchers devised standardised case history forms (Figure 14). These forms performed two primary functions. First, they prompted physicians to collect medical information that was useful to the operation of the computer programme. More importantly, these rigidly standardised forms helped translate the patient's medical picture into numerical information that could be input into the computer. The forms were created, then, to serve as a conduit between two different kinds of media: the material world of pen-and-paper that physicians had been accustomed to for centuries, and the new digital medium of the computer that the Rockefeller team was venturing into for the first time. As Zworykin and Berkley elaborated in an unpublished paper, these forms were created 'in an attempt to arrive at some standardization from which data could be more readily entered into a machine storage'.[330] 'For a computer to accept data', Zworykin reiterated elsewhere, 'it is essential that all of the data supplied

[329] Application for Research Grant, Department of Health, Education, and Welfare, 'Decision-Making Models for Diagnosis and Therapy', 23 Sep. 1968, REP/5/3.
[330] Vladimir Zworykin and Carl Berkley, 'Applications of Medical Data Processing', Paper No. VMe-3-57, VZP/77/69.

to it must be consistent'.[331] Traditional case histories and patient records, heterogeneous and discursively rich, would not suffice. Medical information in such a traditional format was ill fit for transcription onto magnetic tape.

As a longer-term goal, the team imagined that the computer programme would be dynamic, updating disease descriptions on the basis of new and mounting amounts of input data. Jordan Baruch of Bolt, Beranek and Newman (BBN) broached this point in the discussion following a paper presented by Lipkin at the Rockefeller Institute's 1959 conference on 'Diagnostic Data Processing'.[332] Baruch raised 'the possibility of introducing a feedback system whereby the machine improves its performance as time goes on'.[333] Lipkin responded affirmatively. 'We believe a long-term study involving feedback mechanisms should be undertaken. For example, in describing diseases, instead of using the various opinions regarding the significance of the data, or even literature, case hospital data themselves might be used as a basis of the information'.[334] Lipkin thus saw a clear path toward a future where computers were beginning to define disease descriptions 'for themselves'—not unlike the symptom complexes as defined by the MDS.

One member of the hematology group, B. J. Davis, worked with IBM's Taffe Tanimoto to create a computerised programme that could tackle 'the taxonomy problem' in medicine.[335] Tanimoto developed a computer programme that, through application of a mathematical model, could classify cases into disease-like categories.

---

[331] Vladimir Zworykin, 'The Future of Data Processing and World Health', 8 Jun. 1961, draft of presentation at the meeting of the Panel on the Application of Physics to Biology and Medicine, Washington D.C., 14 Jun. 1961, VZP/78/34.
[332] This was the first of what would become a series of IBM Medical Symposia.
[333] Discussion of Lipkin, 'Correlation of Data with a Digital Computer in the Differential Diagnosis of Hematological Diseases', p. 246.
[334] Ibid.
[335] Taffee Tanimoto, 'IBM Type 704 Medical Diagnosis Program', *IRE Transactions on Medical Electronics* ME-7 (1960): 280-283, p. 282.

Essentially, the programme defined certain 'centroid' points of naturally occurring clusters, which purportedly represented the case that is most representative of all other cases in that cluster. The method, when presented at the 1959 Rockefeller conference, stirred considerable discussion. Ledley, while speculating that the method 'would probably be the basis of a good method for finding new diseases which have not been previously defined', also probed, 'Why should the centroid point be the diagnosis? That is not clear in my mind'.[336] Carl Berkeley, another member of the haematology group, responded on Tanimoto's behalf, noting that the clusters should be thought of not as narrow diagnoses but broader classes of illness (e.g., arteriosclerosis versus diffuse degenerative-circulatory disease). Not all attendees were so optimistic about this programme's promise. One physician questioned whether the method would really work in medicine, a domain where clinical pictures frequently overlap and blur together. 'I don't have quite as much faith in medicine's being so well ordered as some here have shown', he warned.[337]

CLINICAL JUDGEMENT, THE MANY-HEADED HYDRA

For a first attempt, the haematology programme performed reasonably well. After examining forty-nine hospital cases, the programme identified the correct disease (as diagnosed by hospital staff)—either as the definitive 'positive diagnosis' or among the shortlist of possible pathologies—in all but two of the cases.[338] Although the

[336] Ibid., p. 283.
[337] Ibid.
[338] For other early work by the hematology team, see R. Ebald and R. Lane, 'Digital Computers and Medical Logic', *IRE Transactions of Medical Electronics* 7 (1960): 283-288; Martin Lipkin, 'Correlation of Data with a Digital Computer in the Differential Diagnosis of Hematologic Diseases', *IRE Transactions of Medical Electronics* 7 (1960): 243-246; Martin Lipkin, Ralph Engle, Jr., B. J. Davis, Vladimir Zworykin, R. Ebald, M. Sendrow, and Carl Berkley, 'Digital Computer as Aid to Differential Diagnosis: Use in Hematologic Diseases', *Archives of Internal Medicine*

computer programme seemed to demonstrate the feasibility of computerised diagnosis, the researchers conceded that it was limited in important ways: the programme considered only a sliver of all haematologic diseases, it omitted complex cases, and it incorporated rough human estimates. This original program, Carl Berkley wrote to Louis Gerstman in 1966, 'reflected the experience of physicians in a relatively cursory way'.[339] In updating the programme, the team's primary aim was twofold. One, they wanted to include additional haematologic diseases and patient findings, and two, they wanted to determine the symptom weights through objective data rather than through the subjective clinical judgment of a panel of haematologists. Lacking access to enough clinical data to determine these symptom weights, the researchers took to systematically reviewing the medical literature. By extracting and combining the data from enough studies and published case histories, they hoped to garner empirically robust information about the incidence of each clinical finding in each haematologic disease of interest.

Led by one of the group's haematologists, M.A. Atamer, this process of culling data from the medical literature to produce a database proved to be 'a tremendous

---

108, no. 1 (1961): 56-72; Martin Lipkin, 'Digital and Analogue Computer Methods Combined to Aid in the Differential Diagnosis of Hematological Diseases', *Circulation Research* 11 (1962): 607-613; Martin Lipkin and Max Woodbury, 'Coding of Medical Case History Data for Computer Analysis', *Communications of the ACM* 5, no. 10 (1962): 532-534; Martin Lipkin and Max Woodbury, 'Analytical Studies Related to the Differential Diagnosis of Hematologic Diseases', *Blut* 9, no. 7 (1963): 449-454; Ralph Engle, Jr. and B. J. Davis, 'Medical Diagnosis: Present, Past, and Future I. Present Concepts of the Meaning and Limitations of Medical Diagnosis', *Archives of Internal Medicine* 112, no. 4 (1963): 512-519; Vladimir Zworykin, 'A Mechanized Matching Procedure for Computer Aided Differential Diagnosis', *Medical Electronics and Biological Engineering* 1, no. 1 (1963): 85-89; Martin Lipkin, Ralph L. Engle, Jr., Betty Flehinger, Louis Gerstman, and M. A. Atamer, 'Computer-Aided Differential Diagnosis of Hematologic Diseases', *Annals of the New York Academy of Sciences* 161 (1969): 670-679.

[339] Carl Berkely to Louis Gerstman, 9 Mar. 1966, REP/14/2.

undertaking'.[340] When Atamer and others turned to clinical data and the medical literature, they encountered what one member of the haematology team characterised as a state of information anarchy. The metaphor captured the situation aptly: physicians collected and reported clinical data according to their own preferences, particularities, and proclivities. A lawless environment bred various problems. The team found, for example, that different physicians referred to different diseases by different names. Von Willebrand disease, a form of haemophilia, went by fifteen different names at the time, and myeloid metaplasia had twenty-five different possibilities in the existing terminology. Atamer also ran into problems arising from a marked degree of variation in how physicians and researchers reported clinical data. For example, if a study didn't report information about a given clinical finding, did that mean that the finding was found to be absent, or simply that it wasn't tested for? On top of this, the Rockefeller team had to grapple with vague and imprecise phrases strewn throughout the medical literature. What did it mean, numerically, when a study stated that a certain symptom or clinical finding was 'infrequent' or 'uncommon' in a given patient population?

As a result of the indeterminacy and heterogeneity of clinical data, the process of collecting and collating data from multiple sources—that is, the process of amassing a database—was not an objective one. It was shot through with subjective judgments.[341] Members of the Rockefeller team had to judge for themselves which studies to include in their review, what to make of studies with anomalous findings, and

---

[340] Progress Report: Electronic Data Processing in Hematology, Grant No. AM-06857-03, U.S. Public Health Services, National Institutes of Health, January 1963, p. 2, REP/14/11.

[341] Lisa Gitelman (ed.), *"Raw Data" Is an Oxymoron* (Cambridge, MA: MIT Press, 2013). Many of the questions Atamer faced (What counts as good data? How much data is enough? What is the optimal way of converting data from one format to another?) are exemplary of what the historian Paul Edwards calls 'data friction'. Edwards, *A Vast Machine.*

how to quantify ambiguous descriptors like 'few' and 'many'. Subjective judgment, they discovered, was recalcitrant; it could not easily be done away with. They attempted to reduce it at one location in the engineering process, only to find it reappear elsewhere. Clinical judgement operated as something of a many-headed hydra—cut one head off, only to find another, or another two, reappear. The need for clinical judgement remained.

The researchers were internally cognisant of this challenge. In a 1962 letter to Zworykin, Berkley described how contemporaneous researchers who were attempting to make sense of data reported from drug tests faced 'the same kind of bottleneck that we have had on the hematologic diagnosis, that is, it requires experts in Biochemistry and Pharmacology to translate from the drug report to the computer format'.[342] An internal progress report on the project, produced in 1967, echoed these sentiments. Pulling data from the world literature, it noted, 'necessitated employing some subjective judgment since different investigators differed in their definitions of certain symptoms as well as in the thoroughness of their reports. One author, for example, might cite 100% incidence of a particular symptom where another author did not cite the symptom at all'.[343]

The persistence of human judgement in increasingly 'rationalised' domains of medicine is a dynamic that runs across the history of medicine. As Harry Marks has shown, the advent and rise of the randomised controlled trial (RCT)—the 'gold standard' of medical research—often failed to permanently settle open questions of therapeutic efficacy and clinical decision making.[344] In fact, RCTs introduced new

[342] Berkley to Zworykin, October 22, 1962, VZP 81/19.
[343] Louis J. Gerstman, Final Report: Electronic Data Processing in Hematology (AM-06857-05): 1 January 1966 - 28 February 1967, p. 4, REP/14/13.
[344] Marks, *The Progress of Experiment.* Also on the limits and challenges of RCTs, see Laura E. Bothwell, Jeremy A. Greene, Scott H. Podolsky, and David S. Jones,

forms of judgement and decision making into the equation. Was the trial adequately designed? Is the study population generalisable? Did researchers adequately adhere to study protocol during the process of implementation? Even after a given RCT produced some result about comparative efficacy, the matter was sometimes far from settled: clinicians continued to debate—and make human judgements about—the quality of the study, the validity of the result, and the extent to which the study's findings applied to *their* particular patients. Clean statistical designs do not supplant the messy political and professional debates around therapeutic choices and 'best practices'.

The University Group Diabetes Program (UGDP) study offers a case in point.[345] After longstanding controversies about the clinical management of diabetes, the UGDP was intended to settle these affairs through state-of-the-art clinical trial design. Contrary to expectations, however, the results of the trial fuelled, rather than squelched, debates about the use of long-term hypoglycaemic drugs in the treatment of diabetes. When the study produced unexpected findings (that the popular hypoglycaemic drug tolbutamide actually *increased* cardiovascular complications), many physicians were forced to decide whether the force of the clinical trial evidence warranted an overriding of what Yale's Alvan Feinstein called the clinician's 'biologic logic' and 'clinical judgement'.[346] In a similar vein, ceding the task of diagnosis to database or machine introduced forms of judgement at both the development end and the implementation end. Developers like Atamer had to employ subjective judgements in the process of

---

'Assessing the Gold Standard—Lessons from the History of RCTs', *NEJM* 374 (2016): 2175-2181; Angus Deaton and Nancy Cartwright, 'Understanding and Misunderstanding Randomized Controlled Trials', *Social Science & Medicine* 210 (2018): 2-21.

[345] Marks, *The Progress of Experiment*, pp. 197-228.

[346] Ibid., p. 213.

creating the database, and clinician-users would have to make subjective judgments about the database's accuracy and fidelity.

TO ERR IS HUMAN?

Databases, because they were judgement-laden, could also be biased.[347] They could reflect systematic errors, lapses, assumptions, and oversights that affect the functioning of a computerised diagnostic system. This possibility did not escape actors at the time. Many framed this problem as one of 'garbage in, garbage out' (GIGO): computer-generated results are only as high of quality as the information entered into them. Biased inputs begot biased outputs.[348] A 1967 article in *Medical World News*, for example, pointed out that computers could only be as rational as their creators.

> Diagnosis requires judgment, and a computer must rely on its human programmers to supply that judgment. The "electronic brain" is no brain at all. It has been called "an idiot machine, capable of generating errors at the same fabulous rate at which it generates correct answers." Program it wrongly and it spews out nonsense. In cybernetic jargon, this awkward situation is called GIGO—"Garbage In, Garbage, Out"—meaning that a computer is no better than the people who run it.[349]

---

[347] Algorithmic bias is a large and lively area of research in STS and the history of technology and computing. See, e.g., Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: New York University Press, 2018); Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Crown Publishers, 2016).

[348] More generally, many developers understood that their systems were limited by data quality. Warner and his colleagued observed, 'Our experience with this series of cases in computer diagnosis indicates that the computer's performance at this time is limited by two factors: (1) the accuracy of the input data from the patient as supplied by the examining physician and (2) the accuracy of the data matrix containing the coincidence of symptoms and diseases'. Toronto, Veasy, and Warner, 'Evaluation of a Computer Program for Diagnosis of Congenital Heart Disease', p. 373. The team likewise noted that incidence figures derived from one patient might not apply to a different population. The cautioned, 'The a priori incidence figures used are specific for this population and would not necessarily apply to patients with congential heart disease referred to another laboratory or clinic'. Ibid., p. 362.

[349] 'Medicine Faces the Computer Revolution: Electronic "Brains" Are Heralding a New Epoch of Improved Diagnosis, Timelier Treatment, and Far Less Medical Paper Work', *Medical World News*, 14 Jul. 1967.

Computers, some at the time realised, did not transcend the human faults of its makers. It incorporated, and even amplified, these faults.

Brodman, too, considered the possibility of biased data. Like the early haematology project, the MDS was 'data-driven'. Brodman and his colleagues did not decide, a priori, how symptoms were weighted, or how symptom complexes were defined. Rather, these weightings and complexes emerged from patterns in the CMI and MDI data. Introduction of bias and error into the data—whether during its creation, processing, or application—could thus translate into inaccurate results. In his 1960 paper 'Diagnostic Decisions by Machine', Brodman considered the analogy between 'the machine method' and the human mind and noted that '[t]he machine method was devised to simulate what is postulated to be the operation of a mind free from any intellectual or emotional bias, and because every mind has many such biases the machine at best can operate analogously only to any isolated human function'.[350] Brodman was careful to note, however, that despite the computer's absent emotionalism, biases could also be introduced into the machine method. 'Biases similar to those associated with decision function can be introduced in a machine system', he warned. This bias entered the system through the back door: through its 'reference data':

> We now have under study machine systems which have been biased in a systematic manner by being supplied with invalid reference data. (This is similar to what is found in people who use invalid information because of ignorance or an emotional disturbance.) Decisions made by these systems are often incorrect in a predictable manner, not because the data are processed illogically but because the logical operations make use of reference data different from those which yield correct decisions. Other systems which have been biased in a random manner (as are individuals with certain central nervous

[350] Brodman, 'Diagnostic Decisions by Machine', p. 219.

> system disorders, such as disseminated atherosclerosis) yield incorrect but logical diagnostic decisions that are unpredictable.[351]

The processing of biased data—no matter how logical or unbiased—generated bad and biased decisions.

Among the most prominent voices arguing against uncritical reliance on data was Alvan Feinstein—the same Feinstein who warned against naïve adoption of clinical trial results.[352] A leading epidemiologist, Feinstein's critiques did not come from a place of nostalgia or mysticism about clinical reasoning.[353] Indeed, in his 1967 book *Clinical Judgement*, called 'revolutionary' by reviewers at the time, Feinstein urged clinicians to adopt 'improved scientific methods for observing and analysing the intact body, mind, existence, and clinical treatment of many'—to make the act of clinical judgement, despite its unavoidable subjective elements, into a kind of 'basic science' of its own.[354] But it was important, Feinstein believed, that physicians bring this clinical judgement—which could be learned, dissected, studied, communicated, and cultivated—to bear in their engagements with clinical data and computers. Feinstein therefore worried about computerised systems based on faulty data—about the consequences of clinically-oriented physicians ceding the realm of diagnosis to professionals with little clinical grounding. As he explained in 1966, 'The right kinds of clinical data and numbers must go into the computer. The big danger is that clinical

---

[351] Ibid., p. 219.
[352] Alvan Feinstein, *Clinical Judgement* (Baltimore: Williams and Wilkins, 1967).
[353] For a nice review of Feinstein's ideas about epidemiology during this period, see, e.g., Alvan Feinstein, 'Clinical epidemiology I: The population experiments of nature and of man in human illness', *Annals of Internal Medicine* 69 (1968): 807-20.
[354] Feinstein, *Clinical Judgment,* p. 380. An underlying point of the book was this: there is nothing mystical about clinical judgement—it can be studied, formalised, and improved.

physicians will stand off and let the mathematicians, Ph.D.s and laboratory-oriented M.D.s take charge'.[355]

Yet, despite glimmers of self-awareness about the problems of bias, judgement, and error, many of Zworykin's contemporaries viewed the computer as relatively immune to the biases that commonly befell human physicians. Humans tire, forget, and get emotionally invested in their work. Computers, by contrast, have no emotions; they don't have bad days. Each case is treated the same as the last. These views were expressed by Edmund McTernan and Dean Crocker in a 1969 issue of *Hospital Physician*. 'No matter how many physicians and related personnel we train, we can't keep up with the demand for health services', they wrote. 'One big reason is that we are *too* human. Our receptors of information are subject to error: Hearing cannot be perfectly trained; vision may fail by degrees, unnoticed; touch is imperfect; personal worries, prejudices, or mind-sets may distort our perception'. They concluded, 'Thus, our ability to gather all the data we need is limited by our humanity, affected by our need to eat, sleep, love, and eventually die'.[356] Their point was simple: to err is human. But their point also had an important corollary: if error was human, accuracy was mechanical.

Such views about the limits of the physician's mind were buoyed by mounting awareness of human-generated medical error through the mid-twentieth century. By the 1960s, and increasingly through the 1970s, numerous studies began to cast aspersions upon the credibility of physicians and the soundness of their decisions. Certain studies, for example, started to reveal a marked degree of regional variation in medical care.

---

[355] Quoted in Harry T. Paxon, 'The Computer: A Report on How It's Affecting the Hospital Physicians', *Hospital Physician*, Sep. 1966, p. 37.
[356] Edmund McTernan and Dean Crocker, 'Push-button medicine is no pipe dream!', *Hospital Physician*, Jan. 1969, p. 85.

One 1969 study showed that, within a single U.S. state and under a single insurance scheme, certain surgical procedures were performed in some regions at a rate three to four times that of other regions.[357] Another study, published four years later in the journal *Science*, documented 'wide variations' in healthcare expenditure, utilisation, and provision even among neighbouring cities.[358] Viewed from a global perspective, these variations were often more stark. In 1968, an international team of researchers from the United States, the United Kingdom, and Sweden found up to fourfold variations in hospital discharge rates and surgical operation rates across the three locales.[359] That a given patient might go under the knife in one zipcode but not in a neighbouring zipcode seemed arbitrary. The existence of regional discrepancies therefore eroded the credibility of physicians and medical institutions. As one commentator would later write, such findings cast 'doubt on [physicians'] knowledge base, on their competence to interpret it, or both'.[360]

---

[357] Charles E. Lewis, 'Variations in the Incidence of Surgery', *NEJM* 281 (1969): 880-884.

[358] John Wennberg and Alan Gittelsohn, 'Small Area Variations in Health Care Delivery', *Science* 182, no. 4117 (1973): 1102-1108.

[359] R. John Pearson, Bjorn Smedby, Ragnard Berfenstam, Robert Logan, Alex Burgess, Jr., and Osler Peterson, 'Hospital Caseloads in Liverpool, New England, and Uppsala: An International Comparison', *Lancet* 292, no. 7567 (1968): 559-66.

[360] Quoted in David Jones, *Broken Hearts: The Tangled History of Cardiac Care* (Baltimore: Johns Hopkins University Press, 2012). For a historical perspective on variation (geographic, racial, and otherwise) in medicine, see chapter 17, 'Puzzles and Prospects', pp. 203-228. Alongside mounting speculation that medical care was inconsistent and ineffectual were anxieties that medical care might actually be injurious. Some of the first systematic studies of clinical error appeared in the 1960s. One of these studies, entitled 'The Hazards of Hospitalization', found that in one hospital 'deleterious episodes befell 20 per cent of all patients admitted to the service' and implored practitioners to remain 'prepared to alter the procedures when imminent or actual harm threatens to obliterate their good'. Elihu M. Schimmel, 'The Hazards of Hospitalization', *Annals of Internal Medicine* 60, no. 1 (1964): 100-110, pp. 108-09. Other formal studies of iatrogenesis that appeared during this time period include O.C. Philips et al., 'The Baltimore Anesthesia Study Committee: Review of 1,024 postoperative deaths', *JAMA* 174 (1960): 2015-84; Max Schapira, Edith Kepes, and Elliott Hurwitt, 'An analysis of deaths in the operating room and within 24 hours of

Proponents of computerisation seized on these findings. A duo of physicians wrote an article in 1966 addressing the question, 'Why use computer as an aid to diagnosis?' The first answer they offered: to reduce errors. Human physicians, they observed, sometimes dropped the ball, whether due to oversight or misdirection:

> Whenever there is a *decision* to be made on diagnosis, prognosis, or choice of therapy, the physician faces the possibility of making the "wrong decision" in relation to the future course of his patient. A number of investigators have reported the intrapersonal and interpersonal *inconsistency* of medical observation and diagnosis. Though inconsistency may not necessarily imply error, physicians would generally agree that significant errors do occur in the practice of medicine. Important errors of decision may be due to *errors of omission*—failing to observer certain facts or failing to consider certain implications—or may be due to *errors of commission*—choosing the wrong course of action, possibly because of a mistaken diagnosis and/or prognosis, or possibly because of error in judgement.[361]

Incorporating computers into the decision making process represented to many one promising way of rooting out medical error.

NUMBERS OVER NARRATIVE

Recognition of the subjective aspects of this process did little to deter the haematology team from pushing forward: after two years of work, Atamer had compiled data from some 18,000 cases reported in the published literature. Atamer arranged this data into a matrix with 76 columns and 538 rows—each column representing a haematologic disease and each row representing a given patient finding. Each cell in the matrix reported the incidence of the particular patient finding in the particular disease. In 1963, Atamer published the information contained within this

---

surgery', *Anesthesia and Analgesia* 39 (1960): 149-157; Robert Moser, *Diseases of Medical Progress*, 3rd edn. (Springfield: C.C. Thomas, 1969).

[361] Stephen R. Yarnall and Richard A. Kronmal, 'Computer aids to medical diagnosis—problems and progress', in *Proc. Conf. Assoc. Comp. Med.* (New York: ACM, 1966), 269-274.

matrix in the form of a book, *Blood Diseases* (Figure 15).[362] Even before the team began to produce computerised disease descriptions from this data, the process of methodically reviewing the world literature generated new ideas about different disease categories in the haematologic field: 'In doing this [creating a database]', the team observed in an annual report, 'it has been found possible to reduce the number of presently recognized disease entities and to combine results from a number of diagnostic tests'.[363]

The information contained in Atamer's matrix formed the basis for the disease descriptions that were input into the computer's memory. As with the original program, the representation of disease in this updated version looked nothing like conventional conceptions of disease. Instead, disease here was essentially a constellation of statistical averages, derived from the population data that Atamer had culled from the literature. Consider the condition of iron deficiency anaemia. In the computer's memory, iron deficiency anaemia was defined as a collection of some 29 clinical variables that were found to be associated, whether positively or negatively, with the disease designation (Figure 16). The strength of the association was expressed in the form of incidence data. Such numerical definitions of disease were a solution to the problem of getting disease concepts into the digital medium of the computer. Their project, Zworykin wrote, 'demonstrated that it is possible to reduce the symptoms of a particular disease into numerical data which can be stored in the memory of a computer'.[364] The

[362] M. A. Atamer, *Blood Diseases* (New York: Grune & Stratton, 1963).
[363] 'The Rockefeller Institute Annual Report', 1 Jul. 1961 to September 1, 1962, Laboratory of V. K. Zworykin, VZP/81/37.
[364] Vladimir Zworykin, 'The Need for Electronics in Medicine', n.d., VZP/82/16.

haematology team hoped that by creating numerical definitions of diseases outside of hematologic field the programme could be easily applied to other fields of medicine.[365]

That their project in computerised diagnosis resulted in numerical definitions of disease was a point of pride among some members of the team. When asked, in 1963, to identify the primary achievements of their work, Zworykin, Engle, Woodbury, and Berkley responded in a joint letter that one of their two chief accomplishments had been 'the numerical specification of the characteristics of diseases to a greater extent and more quantitatively than has previously been attempted'.[366] Zworykin believed that these numerical disease descriptions existed in stark contrast to the discursive descriptions that had dominated much of clinical practice and the medical literature.[367] In a 1961 speech, he derided some disease descriptions from Leavell and Thorup's 1960 textbook *Fundamentals of Clinical Hematology* as 'grossly non-numerical'. To illustrate with a specific example, Zworykin invoked the book's description of agranulocytosis. In a typed draft of the speech, he underlined the portions of the disease description that he believed to be lamentably 'vague': 'Prodromal symptoms such as fatigue, headache, weakness, insomnia, and restlessness usually occur. The majority of patients complain of a sore throat. This often progresses rapidly as chills, increasing fever, and dysphagia develop'.[368] Many physicians may not have found fault with such a description. But compared to disease definitions based on numerical weights or

---

[365] Vladimir Zworykin and Carl Berkley, 'Applications of Medical Data Processing', Paper No. VMe-3-57, VZP/77/69.
[366] Vladimir Zworykin, Ralph Engle, Max Woodbury, and Carl Berkley to Howard Tompkins, 11 Feb. 1963, VZP/82/32.
[367] Vladimir Zworykin, 'The Future of Data Processing and World Health', 8 Jun. 1961, draft of presentation at the meeting of the Panel on the Application of Physics to Biology and Medicine, Washington D.C., 14 Jun. 1961, VZP/78/34.
[368] Quoted in Zworykin, 'The Future of Data Processing and World Health', VZP/78/34. Originally from Byrd Leavell and Oscar Thorup, Jr., *Fundamentals of Clinical Hematology* (Philadelphia and London: W. B. Saunders Company, 1960), p. 321.

statistical averages, individual case histories and discursive descriptions such as these looked nebulous and imprecise. For the haematology team, the project of computerising diagnosis was inextricably linked to the project of 'obtain[ing] accurate probabilistic statements to replace vague disease descriptions'.[369]

To be sure, computers did not spawn numerical representations of disease. By the mid-nineteenth century, for example, hospital-based physicians in Great Britain and France began to produce numerical, rather than narrative, disease descriptions of rheumatism.[370] Nevertheless, this early project in computerised diagnosis demonstrates the potential role of the computer as a key driver behind the evolution from clinical descriptions that were discursive to clinical descriptions that were numerical and aggregative. The latter, unlike the former, was suitable for a novel digital medium. In many ways, the spreading use of computers created a new *cognitive model* about what it meant for medical knowledge to be truly standardised.[371] While the content of Leavell and Thorup's text surely looked structured by comparison to discursively rich case studies of the nineteenth century, early pioneers in biomedical computing instead saw chaos and anarchy. They were working with a new and different cognitive model.

---

[369] Vladimir Zworykin, Ralph Engle, Max Woodbury, and Carl Berkley to Howard Tompkins, draft of letter, 1 Feb. 1963, REP/14/2.

[370] Peter English, 'Emergence of Rheumatic Fever in the Nineteenth Century', in Charles Rosenberg and Janet Golden (eds.), *Framing Disease* (New Brunswick: Rutgers University Press, 1992), 20-32.

[371] In this discussion on cognitive models, I am drawing upon John Harley Warner regard of therapeutics as not just social or material practices but a 'cognitive system' as well. See John Harley Warner, *The Therapeutic Perspective: Medical Practice, Knowledge, and Identity in America, 1820-1855* (Princeton: Princeton University Press, 1986). Recently, Warner has extended these observations, discussing the 'cognitive models' that both produced and emerged out of changing ways of organising the medical record through the nineteenth century. John Harley Warner, 'Vital Signs: The Transformation of the Medical Record in the Nineteenth Century', keynote address at 'Media Medica: Medicine and the Challenge of New Medicine', conference at Center for Medical Humanities and Social Medicine, Johns Hopkins University, Baltimore, Maryland, 17 Oct. 2017.

- CHAPTER IV -

The Medical Mind

'The coming of high speed computers has been the wind of change for many of the sciences and their applied technologies'.[372] So John A. Jacquez opened a conference on 'The Diagnostic Process' held at the University of Michigan in 1963. In additional opening remarks, the Dean of University of Michigan's School of Medicine, William N. Hubbard, Jr., further clarified the circumstances leading up to the conference and identified some of its goals:

> We are aware of the interest that has developed in the last few years in the use of mathematical models and electronic equipment in the diagnostic process. We have not been confident, however, that we understood the intrinsic nature of the diagnostic process well enough so that we could proceed with the kind of transitional research that would carry us effectively into use of mathematical models and machine data processing. We have had some concern that one might adapt the diagnostician to the device rather than vice versa, if one did not have clearly in mind the role that the mathematical model and the machines can appropriately play.[373]

With computers knocking at the diagnostician's front door, the conference organisers saw it as an 'opportune time to re-examine some of the features of the diagnostic process'—to corral together a group of experts to discuss and define how diagnostic decisions might be made.

The haematology team was well-represented at the conference. Ralph Engle delivered the opening paper on 'Medical Diagnosis', where he proposed a philosophical framework for thinking about how physicians reason and classify

---

[372] John A. Jacquez, 'Preface' in John A. Jacquez (ed.), *The Diagnostic Process*: *Proceedings of a Conference Held at the University of Michigan*, *May 9-11, 1963* (Apr. 1964), iii-iv, p. iii.
[373] William N. Hubbard, 'Welcome and Prologue', in John A. Jacquez (ed.), *The Diagnostic Process*: *Proceedings of a Conference Held at the University of Michigan, May 9-11, 1963* (Apr. 1964), 3-5, pp. 3-4.

diseases. Martin Lipkin's paper, though more technical, also touched upon the question of how physicians think—and whether the computerised system approximated such human reasoning. As interdisciplinary teams worked to automate medical diagnosis, they repeatedly bumped up against fundamental epistemological questions. How do doctors think? What is the nature of human reasoning? Are human modes of reasoning amenable to computerisation? If they are, *should* computers aim to mimic them? In what ways do physicians come to define and classify disease categories? Are these categories even necessary? While carrying the haematology programme into its final decades, this chapter explores how computers both generated and incorporated ideas about the medical mind and its inner workings.

MEDICINE'S METAPHORS

The creep of computational tools into medicine gave researchers and practitioners a new vocabulary for understanding and framing the fundamental tasks of medicine. Increasingly through the 1960s and 1970s, developers of computerised diagnostic systems began to see medicine through the lens of computational metaphors. To them, medicine seemed to be little more than an information processing discipline. Their immediate material world—the world of punch cards, magnetic tape, processing units, and even yes-no surveys—shaped their cognitive world, their way of understanding medical work. These new cognitive outlooks, in turn, circled back to fuel and justify their material practices and research efforts.

The power of medicine's metaphors—to shape patient experience and drive therapeutic decision making—has been richly described by historians and social critics alike. Susan Sontag famously argued that the metaphors around disease are not just 'mere' matters of language: they generate stigma, shade public perception of the

disease, and modify patients' own experiences of suffering.[374] These metaphors, she contended, 'imply living processes of a particularly…horrid kind'.[375] Whereas Sontag focused on the negative valence of medicine's metaphors (and argued for a stripping down of medicine to its mechanistic core), other historians have shown that many disease metaphors can also perform cultural work with more mixed and nuanced repercussions. Keith Wailoo, for example, has argued that sickle cell anaemia embodied, among other things, 'the unrecognized pain of the African American community' during the American Civil War.[376]

Metaphors can also structure therapeutic decision making and the medical field's engagements with its patients. The historian David Jones, for example, has demonstrated how understandings of the heart and cardiovascular system as a system of 'pipes' and 'plumbing' affected intuitions about which cardiac therapeutics might work—and why. The metaphor that heart attack and angina were little more than problems of clogged pipes (i.e., obstructive atherosclerotic plaques) generated a powerful intuitive logic that surgical approaches of undoing or bypassing the clog would surely remediate, if not cure, the condition. Jones explains, 'The plumbing analogy makes a powerful claim: since angioplasty or surgery can relieve the

---

[374] Susan Sontag, *Illness as Metaphor* (New York: Farrar, Straus & Giroux, 1978); Susan Sontag *AIDS and Its Metaphors* (New York: Farrar, Straus & Giroux, 1989).
[375] Sontag, *AIDS and Its Metaphors*, p. 9.
[376] Keith Wailoo, *Dying in the City of the Blues: Sickle Cell Anemia and the Politics of Race and Health* (Chapel Hill, NC: University of North Carolina Press, 1997), p. 3. Historians and others have countered Sontag's call to expunge medicine of metaphors, arguing that metaphors are an inevitable part of human life and that 'doctors and patients need to communicate through the rich possibilities of metaphor'. Benjamin J. Oldfield and David S. Jones, 'Languages of the Heart: The Biomedical and the Metaphorical in American Fiction', *Perspectives in Biology and Medicine* 57, no. 3 (2014): 424-442, p. 424.

obstruction or bypass it altogether, they must work. The pairing of theory and therapy is so persuasive that excellent outcomes seem inevitable'.[377]

That computers and protocols—like certain cardiac interventions—have ridden on the coattails of certain metaphorical and rhetorical constructions has not escaped historians and STS scholars. In his 1995 article 'Turning a Practice into a Science', Marc Berg charts a shift during the post-war period in the rhetoric around what 'scientific medicine' was.[378] In the early post-war years, Berg argues, clinicians saw the pursuit of scientific medicine as an *external* pursuit, one that required the appropriate management of the extrinsic, socio-economic factors that impinged upon the medical encounter. By this rendering, physicians and others worked to rationalise medicine through reformation of the external obstacles—financing structures, government oversight, and so forth—that stood in the way of 'the artful application of medical science'[379]. Later things changed. Increasingly physicians saw the challenge of rationalising medicine as an *internal* challenge—one defined by, and to be achieved through, the disciplining of the physician's mind. This 'cognitivist' discourse, Berg observes, 'located the scientific status of medical practice in the mind, and also likens the physician's mind to a computer'.[380] Under this discourse, the threats to a rational medicine come from *within* the physician—from the human biases and errors that can get in the way of a mechanical, rote application and interpretation of medical information.

Berg argues that this rhetorical, discursive, and metaphorical shift allowed for the proliferation of decision support tools throughout medical practice. '[O]nce the

---

[377] Jones, *Broken Hearts,* p. 28.
[378] Marc Berg, 'Turning a Practice into a Science: Reconceptualizing Postwar Medical Practice', *Social Studies of Science* 25, no. 3 (1995): 437-476.
[379] Ibid., p. 444.
[380] Ibid., p. 445.

diagnostic task is seen as an exercise in statistics', Berg reasons, 'physicians' performance is seen as second-rate'.[381] What Berg's account misses, however, is the arrow going in the opposite direction: he overlooks the extent to which the introduction of computers into medicine indeed benefitted from, but also *created*, ideas about medicine's computational and informational underpinnings.

As people like Zworykin began working on computerised approaches to diagnosis, they began to see medicine differently. Physicians were computers; patients, repositories from which data could be extracted. Zworykin quoted such a view in a talk on 'The Need for Medical Electronics': 'Thus I see a situation—50 to 100 years in the future—in which a physician will extract the proper information from the patient and dial this information in a central computing agency. The computer will sort its stored funding of knowledge, make the indicated comparisons and calculations, and send back a report'.[382] Others agreed. 'The diagnostic process really is a kind of information retrieval problem…', observed one audience member at a talk delivered by Engle.[383]

Some of Zworykin's contemporaries rethought the entire medical system along these lines. Starting in 1964, IBM undertook the development of the Clinical Decision Support System (CDSS) 'to aid in the acquisition of patient data, and to suggest appropriate decisions on the basis of the data thus acquired'.[384] The CDSS broke down the decision making process into microdecisions and used threshold logic to generate

---

[381] Berg, *Rationalizing Medical Work*, p. 45.
[382] Zworykin, 'The Need for Electronics in Medicine', p. 43 VZP/82/16.
[383] Engle, 'Medical Diagnosis', from *The Diagnostic Process: Proceedings of Conference at University of Michigan, May 9-11, 1963*. Jacquez (ed.) (Ann Arbor, Michigan: 1964), REP/10/ 'Speech, Medical Diagnosis – The Diagnostic Process, Proceedings of Conference at Univ. Michigan, May 9-11, 1963'.
[384] Gerald Goertzel, 'Clinical Decision Support System', Talk given at the New York Academy of Sciences' Conference on the Use of Data Mechanization and Computer in Clinical Medicine, 16 Jan. 1968. REP/3/'Clinical Decision Support System, 1967-1972'.

decisions.[385] Overseeing this effort was the physician-turned-engineer Frederick J. Moore.[386] In June 1968, Moore published a report on the effort, analysing the medical system's defining challenges and arguing for their computerised decision support system as a fitting solution. The way Moore argued for this fit was through a reframing of the medical system as an information system. This reframing, Moore claimed, was less of an interpretation than a concrete 'finding' of the report: 'In this report we examine medical care as a system, find that it is primarily an information system in which the decision element is crucial…'.[387] The report reformulated nearly every aspect of medicine—from the tectonics of the system to the granular details of the clinical encounter—as matters of information processing:

> The function elements of health maintenance … are summarized as follows:
> 1. The acquisition of data to establish health status,
> 2. The comparison of these data with stored knowledge,
> 3. The decision as to
>   a. what added data are needed
>   b. what treatment is indicated,
> 4. The administration of the treatment and the acquisition of follow-up data, and
> 5. The addition of all this experience to stored knowledge.[388]

Having outlined the medical system's core elements, Moore felt comfortable concluding that medicine's defining processes, opportunities, and challenges were informational in nature: 'Thus, it appears that the medical care system is predominantly informational—i.e., it depends largely on the acquisition, storage, and interpretation of information by both the patient and the doctor—and that the decision process is the pivotal element in the system, toward which all data acquisition is directed and upon

---

[385] Ultimately the CDSS failed in large part because physicians had trouble building the data base on which the system depended.
[386] 'Frederick J. Moore, IBM Diagnostician', *New York Times*, 15 Oct. 1972, p. 72.
[387] Frederick J. Moore, *Development of a Clinical Decision Support System* (Yorktown Heights, NY: IBM, 1968), p. 1. REP/3/'Clinical Decision Support System, 1967-1972'.
[388] Ibid., pp. 3-4.

which the action taken by the patient or by the doctor depends'.[389] Such a view recast the role of patients and physicians as sources and consumers of data: 'The doctor, as well as the patient, stores, retrieves, and processes data regarding the patient's present disorder and data from relevant clinical experience'.[390] Viewed as an information processor, the doctor's shortcomings came to the fore. The doctor, Moore claimed, relies upon 'two forms of information storage: human memory and records'.[391] The latter, however, is notoriously error-prone: 'it is subject to bias, false drops, and failures, a major factor governing retrieval being the area of interest or attention at the moment'. If medicine was all about data—its creation, storage, management, processing, and distribution—the computer held great promise.

This informational view of medicine gave perceptions of a particular crisis in medicine—that of the 'information overload'—sharper outlines and greater urgency. Many physicians worried that the medical field was facing a threat that was almost Malthusian in nature: whereas human cognitive abilities remained relatively fixed over time, the quantity of medical knowledge available to the clinician seemed to be growing exponentially. The result? Too much information for even the academic medical specialist, let alone the general practitioner, to know.[392] If medicine was, fundamentally, about the accumulation, management, and application of 'information', as many increasingly believed, then concerns about information chaos and overload

[389] Ibid., p. 4.
[390] Ibid., p. 5.
[391] Ibid., pp. 5-6.
[392] Concerns about the threat of information overload were nothing new, as Anne Blair has shown. These anxieties were felt particularly, though, among twentieth-century physicians, who saw a host of biomedical breakthroughs. Anne Blair, *Too Much to Know: Managing Scholarly Information before the Modern Age* (New Haven: Yale University Press, 2010).

presented real threats. The haematology team, like so many others at the time, isolated this as a driving motivation of their program.

Most medical and scientific commentators spoke of this phenomenon in rather general terms, gesturing to 'the exponential explosion of knowledge' unfolding around them.[393] But one researcher, David Durack of Duke University, adopted a more quantitative, if peculiar, approach. In his classic (and jocular) 1978 article in the *New England Journal of Medicine*, 'The Weight of Medical Knowledge', Durack sought to 'take some objective measure' of perceptions among physicians concerning the impossibility of keeping abreast of the current literature.[394] Durack's 'objective measure' of choice was the *Index Medicus*, a comprehensive bibliographic index of all journal articles published in the medical sciences. Durack found that the *Index*, as measured by weight, was indeed undergoing exponential growth: between 1879 and 1945 the *Index* hovered steadily around two kilograms, before doubling in weight between 1946 and 1955, and increasing another seven-fold, to reach nearly thirty kilograms, between 1955 and 1978 (Figure 17). It therefore comes as no surprise, Durack quipped, that '[m]ost of us feel weighed down by the heavy and increasing burden of medical reading'.[395] Despite its playful tenor, the article gave voice to real and pervasive anxieties about the challenge of providing quality, up-to-date medical care in an era of information overload.

---

[393] Bryan Williams, 'Personal View', *BMJ* 2, no. 23 (June 1973), p. 712. See also George Pickering, 'Medicine and Society—Past, Present, and Future', *BMJ* 1, no. 23 (January 1971): 191-6, p. 192. In 1963, Derek de Dolla Price famously charaterised the situation more empirically. Derek J. de Sollar Price, *Little Science, Big Science* (New York and London: Columbia University Press, 1963).
[394] David Durack, 'The Weight of Medical Knowledge', *NEJM* 298, no. 14 (1978): 773-5, p. 773. The article was read widely by contemporaries, with many calling it 'a classic'. See Dr. Diane Madlon-Kay of the St. Paul-Ramsey Medical Center, quoted in 'Science Watch; Heavy Reading', *New York Times*, 3 October 1989.
[395] Durack, 'The Weight of Medical Knowledge', p. 773.

The informational metaphors adopted by those like Moore populated not just the world of language but also the world of image. Renderings of the physician as computer, of the patient body as punched card, of medical knowledge as 'information' swarmed press and popular accounts of computerised diagnosis. In one such image, published in 1968, a traditional anatomical drawing of the human body appears on a punched card, recasting the patient's body as 'data'—as a material input ready for computer processing (Figure 18).[396] Other images portrayed the physician as this computer. One such image, published in *Hospital Physician*, presented the physician as a mosaic, constructed out of wires, gears, nuts, and bolts (Figure 19).

COMPUTERISATION AS VALUATION

Renderings of the physician as a machine tapped into a familiar set of fears. Analogies between mind and machine suggested to many that physicians could be replaced. With some, like Homer Warner, imagining futures where computers take on medicine's cognitive heavy-lifting and turn physicians into mere 'well-trained data-collectors', others pushed back, emphasising the fundamental aspects of clinical work that seemed to defy formalisation and computerisation. The former president of the American Medical Association (AMA) held such a stance, one that surely pleased many members of his former organisation. While discussing the potential benefits of automation to practicing physicians, Orr at the same time reassured, '[t]here is no danger of practices being taken over by computers. Some need for personal contact is what makes the practice of medicine an art as well as a science'.[397]

---

[396] 'Computer and Clinical Medicine', *The Sciences* 8 (1968). REP/3/'Computers in Medicine, 1965-69'

[397] Quoted in 'Biomedical Electronics: A report on an important new bridge connecting mathematics, engineering and medicine, with an emphasis on human systems', n.d., REP/3/'Computers in Medicine, 1959-1964'.

In defending medicine against the threat of automation, many physicians espoused what the historian Christopher Lawrence has called an 'epistemology of incommunicable knowledge'.[398] In his study of the British medical profession from the mid-nineteenth to early twentieth centuries, Lawrence identifies a rhetorical tradition whereby traditional physicians invoked an 'epistemology of individual experience' to 'defend the autonomy of clinical practice'. By describing clinical practice as an art that defied analysis, these physicians cordoned off clinical medicine from encroachment by new technologies like the sphygmomanometer and new 'scientific' modes of analysis.[399]

Although physicians in the mid-twentieth century faced the spectre of new digital technologies, they drew upon these longstanding rhetorical strategies. They emphasised the practice of medicine, and the work of diagnosis, as artful, intuitive, tacit. Excellence in these realms, they argued, derived more from years of irreducible experience rather than from pages of book-bound knowledge. Even developers of computerised systems adopted this language to allay fears about deskilling and replacement. Describing the diagnostic process as indefinite and noncommunicable made both intuitive and strategic sense, as many proponents of computerisation were arguing the flip side of things. One physician, Dr. Iker, reasoned during the 1963 conference on 'The Diagnostic Process': 'I would submit that if clinical diagnosis is a science then a computer can do the job better than a clinician. If it is scientific, it is

---

[398] Christopher Lawrence, 'Incommunicable Knowledge: Science, Technology and the Clinical Art in Britain 1850-1914', *Journal of Contemporary History* 20, no. 4 (1985), 503–20.

[399] Ibid., p. 505. Beyond Lawrence, for a synthetic historiographical discussion of skill and authority in the history of medicine, see Nicholas Whitfield and Thomas Schlich, 'Skills through History', *Medical History* 59, no. 3 (2015): 349-60. For instructive case studies, see Anderson, 'The Reasoning of the Strongest'; David S. Jones, 'Visions of a Cure: Visualization, Clinical Trials, and Controversies in Cardiac Therapeutics, 1968-1998', *Isis* 91, no. 3 (Sep. 2000): 504-41.

communicable and if it is communicable then a computer obviously can do it more quickly and more reliably'.[400]

Questions about how to describe and represent computers nagged at the haematology team as well. Ralph Engle's perspective on the nature of the computer—and its relationship to the human—was rather more nuanced than that of many of his contemporaries. Pioneers in biomedical computing frequently lamented that physicians seemed to suffer from the 'computer allergy' more severely than did any other professional group. The unreceptive attitude of many traditional practitioners was widely recognised as a—if not *the*—central barrier to getting computing technologies out of the research laboratory and into the clinic.[401] In 1966, Stephen Yarnall and Richard Kronmal characterised the 'psychological resistance' among physicians to computers in diagnosis. In their view, this resistance had three primary facets:

> [A]ny workable system for computer-aided diagnosis must consider the psychological resistance among physicians against the use of computers for medical diagnosis. There are three main points in this regard: First, physicians have traditionally approached diagnosis in their own individual manner, and attempts at systematization are generally unwelcome; second, many of the systems introduced to-date have required more, rather than less, work by the physician with no immediate gain apparent to him or his patient; third, physicians, and many laymen, fear computers as strange and impersonal machines which may destroy meaningful doctor-patient relationships and even displace physicians to some extent.[402]

Engle shared these concerns about an unreceptive, even antagonistic, clinical audience. But it was actually the opposite attitude—uncritical acceptance—that came to trouble Engle more deeply. Physicians, he believed, needed a deep understanding of

---

[400] Discussion of Martin Lipkin, 'Data Processing in the Diagnostic Process', in John A. Jacquez (ed.), *The Diagnostic Process*: *Proceedings of a Conference Held at the University of Michigan*, *May 9-11, 1963* (Apr. 1964), 255-280, p. 275.
[401] The resistance among traditional medical practitioners toward new medical technologies is a recurring theme in the history of medicine.
[402] Yarnall and Kronmal, 'Computer aids to medical diagnosis—problems and progress'.

how a computer programme worked; only then would they be able to effectively evaluate its clinical value and, more crucially, its limitations: 'It concerns me that some physicians are only too anxious to let a laboratory test or a computer make their decision. These tools should only be one input into the decision-making process. The physician must add his own perspective'.[403] The output of a computer program, Engle warned, can convey a false decisiveness that physicians may be all too eager to accept uncritically. 'The aura of finality and correctness of the computer', Engle predicted, 'will be difficult for some physicians to overcome'.[404] The single computer output does little to convey the messy, contingent, partial, and uncertain choices that the developers made in creating the programme.

Engle, it seems, feared that a 'living' metaphor about the computer risked gradually shading into a 'dead' one; that is, he worried that metaphors may lose their status as mere figures of speech and begin to be interpreted as real or literal. As a result, when Engle did invoke metaphors, he was careful to point out their status as such. His views on this subject were heavily informed by the work of the German philosopher Hans Vailhinger. In his book *The Philosophy of 'As if'*, Vailhinger argues that much human knowledge is a mere assemblage of 'fictions' but that there is often value in acting 'as if' these fictions were true for pragmatic reasons.[405] In much the same way, Engle believed it a 'fiction' to speak of a computer's 'intelligence', but he nevertheless granted that such language could be justified pragmatically. He reasoned, '[G]reat

---

[403] Ralph L. Engle, 'Implications of Computer-Aided Medical Consulting', n.d., REP/6/5.
[404] Ralph L. Engle, 'Medical Diagnosis: Present, Past, and Future. III. Diagnosis in the Future, Including a Critique on the Use of Electronic Computers as Diagnostic Aids to the Physician', *Archives of Internal Medicine* 112, (1963): 530-543, p. 540.
[405] Hans Vailhinger, *The Philosophy of 'As if': A System of the Theoretical, Practical and Religious Fictions of Mankind*, trans. C. K. Ogden (New York: Harcourt Brace, 1925).

benefits can result from acting "as if" models of artificial intelligence were real life human intelligence especially if it may lead to an effective coupling of the computer with the human brain'.[406]

Other members of the haematology team also often acted and spoke 'as if' the computer were a kind of human—although surely with less of a philosophically-inflected self-awareness about what they were doing. The storage of incidence data in the computer's memory, the team wrote in a 1963 progress report, represented a kind of 'experience' not unlike the physician's own experience: 'This type of information stored in the "permanent" computer memory will eventually constitute a body of "computer experience" which will exceed human capability'.[407] In a letter written the same year, members of the team similarly discussed the value of 'the cumulative experience of a computer'.[408] Many others took these metaphors further, assigning computer systems varying degrees of educational attainment—'medical students', 'residents', 'attendings'—based on the quantity and quality of information stored within them.

As with Brodman's case, such descriptions of the computer as an active participant, one possessing its own body of experience, may speak to how the team experienced the computer's role in making and remaking disease concepts. Through the process of computerising diagnosis, new understandings and definitions of disease came into view. 'When we started this', Zworykin noted in 1961, 'no one knew which

[406] Engle, 'Implications of Computer-Aided Medical Consulting', p. 3, REP 6/5.
[407] 'Progress Report: Electronic Data Processing in Hematology', Grant No. AM-06857-03, U.S. Public Health Services, National Institutes of Health, Jan. 1963, p. 5, REP/14/11.
[408] Vladimir Zworykin, Ralph Engle, Max Woodbury, and Carl Berkley to Howard Tompkins, 11 Feb. 1963, VZP/82/32.

relevant symptoms might be reported in hematologic diseases…'.[409] Once their project was well under way, they had a clear—indeed quantitative—sense of how patient findings related to different haematologic diseases.

Decisions about where computers ought to go—what parts of medicine they were well suited to automate—were not just technical decisions. Behind these decisions were implicit valuations of different kinds of medical work. As Stephanie Dick has pointed out, new digital technologies create 'new economies of intelligence' that dictate 'whose cognitive labor is valued and devalued, displaced or replaced…'.[410] To carve out certain aspects of medicine for computerisation was to partition medical work into tiered categories: work that was human, cognitive, artful (high value), and work that was mechanical, mindless, rote (low value). Many developers of early computerised diagnostic systems insisted that computers would only take over the 'rote' aspects of medicine, leaving physicians able to dedicate more of their time toward the human, and humanistic, dimensions of medicine.[411] Leave to the computer the rote and the routine; leave to the physician the complex and contradictory. The two, mind and machine, were different but complementary. As the pioneer in the computerisation of ECG Cesar Caceres wrote in 1968, 'Proper use of both the machine's tireless ability for routine analysis and the physician's inimitable ability to put together sensed or intuitive impressions is the key to improvements in tomorrow's

[409] Vladimir Zworykin, 'The Future of Data Processing and World Health', 8 Jun. 1961, draft of presentation at the meeting of the Panel on the Application of Physics to Biology and Medicine, Washington D.C., June 14, 1961, VZP/78/34.
[410] Stephanie Dick, 'Artificial Intelligence', *Harvard Data Science Review* (2019). Available at: https://doi.org/10.1162/99608f92.92fe150c
[411] These same arguments have enjoyed a second life in recent years. See, e.g., Eric Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (New York: Basic Books, 2019).

diagnostic practice'.[412] Computers, these physicians and engineers suggested, would take over only the subset of the physician's work that was *beneath* him. In doing so, the physician would be 'free' to pursue the aspects of medicine that he found more interesting and rewarding.

Members of the haematology team thought along similar lines. Engle believed that the diagnostic process was not just one thing but represented a collection of heterogeneous activities. It comprised aspects suited to the computer *and* aspects suited to the human. He thus described the need to 'dissect out' from the diagnostic process those aspects which were routine and rote. Engle explained,

> Human beings … are able, in some unexplained way, to grasp subtle and complex relationships, recognize intricate patterns, recall integrated and interpreted facts, and originate new ideas. Computers have great potential for aiding thought if the mechanical facets of reasoning of identified and delegated to the computer. In order for computers to realize their full potential, the problem of interaction between human beings and computers must be approached from both sides. People must learn to analyse their thought processes so that they can dissect out the mechanical part and formulate the functions the computer is to perform.[413]

Here, it is clear how decisions about how to computerise diagnosis always implied certain ideas about humans and human nature—what kinds of work physicians were capable of, interested in, suited for. As computers entered the arena of medical diagnosis, then, they not only came up against traditional ideas of professional identity; they also constructed new ideas about the physician's identity. Computers forced new discussions and thoughts about the physician's 'appropriate' role in the diagnostic process and the clinical encounter.

---

[412] Cesar Caceres, 'Computer diagnosis, off and running', *Medical World News*, 17 May 1968.
[413] Engle, 'Implications of Computer-Aided Medical Consulting', n.d. p. 3,, REP/6/5.

HOW DOCTORS THINK

Questions about how doctors think—and whether computers ought to replicate such processes—pestered Zworykin's group almost from day one. What does clinical reasoning even look like? (How) can it be distilled, studied, visualised? The haematology team believed they found a way of answering such questions in the Clinicopathological Conference (CPC), a formalised event where physicians reason aloud through a selected patient case. By analysing published CPCs—and by staging their own—Zworykin and his colleagues hoped to distil the process by which the physician used medical information to reach a diagnosis.

The CPC evolved out of a broader late-nineteenth-century movement to introduce case studies into education. The so-called 'case method' of medical education was an import from the legal academy. Under the direction of its visionary dean, Christopher Langdell, Harvard Law School pioneered the case method of legal instruction, whereby law professors, instead of relying on legal theory and abstract principles, would teach different areas of law by thinking through actual legal cases. Across the Charles River, reformers at Harvard Medical School's campus soon realised that an analogous method could be applied to medical education.[414] Medical instructors

[414] In explaining the medical school's adoption of the case method, some accounts have emphasised the role of Walter Cannon, a fourth-year medical student at Harvard Medical School. Cannon's roommate at the time, a second-year student at Harvard Law School, was enthusiastic about the law school's use of specific legal cases in education. Inspired by his roommate's enthusiasm, Cannon advocated, in a 1900 article in the *Boston Medical and Surgical Journal*, for a similar style of teaching in medical education. This is an interesting story but it detracts from the broader cultural and ideological forces at play. See Nancy Lee Harris and Robert Sculley, 'The, Clinicopathological Conferences (CPCs)', David N. Louis and Robert H. Young (eds.) *Keen Minds to Explore the Dark Continents of Disease* (Boston, MA: Massachusetts General Hopistal and Harvard Medical School, 2011), 349-362. On the history of medical education, see Kenneth Ludmerer, *Learning to Heal: The Development of American Medical Education* (New York: Basic Books, 1985); Kenneth Ludmerer, *Time to Heal: American Medical Education from the Turn of the Century to the Era of Managed Care* (New York and Oxford: Oxford University Press, 1999).

increasingly saw pedagogical value in presenting well-selected patient cases as a way of illuminating aspects of disease and diagnosis. It was precisely this logic that motivated Richard C. Cabot, then a visiting physician at the Massachusetts General Hospital (MGH), to begin scheduled performances of the CPC sometime around 1909.[415] At the conference, the featured physician would be presented with descriptions of a patient's symptoms along with laboratory and radiological information. Standing before an audience of physicians and unwitting of the actual diagnosis, the physician would explicate his diagnostic reasoning and come to a diagnosis for the case before him.[416] The CPC was therefore a performative exercise, and it occasionally took on the character of a drama.

The purpose of the CPC was less about making a correct diagnosis than about simulating the diagnostic process itself. The editors of a 1960 volume of CPCs called the conference 'an exercise in deductive reasoning'—one in which '[i]t is less important to pinpoint the correct diagnosis than to present a logical and instructive analysis of the pertinent conditions involved'.[417] The complexity of solving many diagnostic problems justified this performance of the diagnostic process. Not only did physicians face vast amounts of patient information; they had to synthesise medical

---

[415] Christopher Crenner, 'Diagnosis and Authority in the Early Twentieth-Century Medical Practice of Richard C. Cabot', *Bulletin of the History of Medicine* 76 no. 1 (2002) 30-55, p. 37. On Richard Cabot more generally, see Christopher Crenner, *Private Practice in the Early Twentieth-Century Medical Office of Dr. Richard Cabot* (Baltimore: Johns Hopkins University Press, 2005).

[416] In 1924, the *Boston Medical and Surgical Journal*—the precursor to the *New England Journal of Medicine*—started to publish the Clinicopathological Conferences as Case Records of the Massachusetts General Hospital. Richard C. Cabot and Hugh Cabot, eds., 'Ante-mortem and Post-mortem Records as Used in the Weekly Clinico-pathological Exercises. Case 10271', *Boston Medical and Surgical Journal* 191, no. 4 (1924): 30-34.

[417] Benjamin Castleman and H. Robert Dudley (eds.), *Clinicopathological Conferences of the Massachusetts General Hospital, Selected Medical Cases* (Boston: Little, Brown and Co., 1960), p. viii.

information coming from different sources and reconcile findings that seemed to point in different directions. In complex cases, it was rare for a patient's medical picture to map neatly onto any textbook description of disease. The CPC was a demonstration in how to navigate this complexity. Amidst the bustle of medical work going on in the hospital, the CPC purported to isolate—and to stage—clinical cognition in its purest form. For some, the exercise promised a window into the inner workings of the physician's mind.

It was for these reasons that the haematology team took an interest in the CPC.[418] Yet, the team ultimately concluded the CPC offered them little actionable insight. Engle later explained, 'Obvious approaches to answer some of the questions concerning how a physician makes a diagnosis include analyses of published Clinical Pathological Conferences (CPC) such as those reported in the *New England Journal of Medicine* or the *American Journal of Medicine*, and interviews of physicians as they go through the diagnostic process during a real or simulated Clinical Pathological Conference. When our group first became interested in computer aids to decision-making in the 1956 we used both of these approaches and came to a few conclusions'.[419] Specifically, Engle and his colleagues realised that the conditions around CPCs were simply too sanitised and unrealistic to be of much use: 'CPC's are carefully selected situations and are not representative of patients in general. They are more like games than like real life situations since the physicians know that the patients have died and that the decision reached will not have any effect on the treatment of the

[418] Others also viewed CPCs as 'excellent windows through which to study the diagnostic process'. David Eddy and Charles Clanton, 'The Art of Diagnosis: Solving the Clinicopathological Exercise', *NEJM* 306, no. 21 (1982): 1263-1268, p. 1267
[419] Ralph Engle and Betty Flehinger, 'The Diagnosis Process I. Role of Physician', unpublished manuscript, 1978, p. 14, REP/11/1.

patients'.[420] The exercise, to be fair, was not wholly fruitless. The CPCs did, Engle acknowledged, 'give some indication as to how a physician approaches a particular problem'—namely, physicians 'often pick from a multitude of findings certain ones of high significance and proceed to do a differential diagnosis related to those findings'.[421] (This was the conclusion that Brodman had also reached in his own, more inwardly reflective, analysis of how doctors think.)

Ultimately, though, the haematology team found that physicians were too often unable to explicate their decision making processes. They could effectively navigate through some diagnostic maze, but they were not particularly good at describing *how* they got through—why, for instance, they made *this* assumption, and not *that* assumption. Different physicians, moreover, sometimes took different paths through the diagnostic process, even if they arrived at the same diagnosis. The team was left wondering: was one of these routes better, or more representative, than the other? As Engle elaborated, 'The physician frequently could not give a rational reason for making certain inferences and certain decision even though he may have been satisfied with his line of reasoning. There was considerable variation in approach from one physician to another and by the same physician in analysing different cases. The physician often thought in the most general terms without much specific information. Occasionally, the basic information was erroneous. There did not seem to be anything specific about medical reasoning that was not similar to decision-making outside the field of medicine'.[422]

To many it was not self-evident that researchers should *want* to replicate human reasoning, even if they could. Computers, some believed, brought a different set of

[420] Ibid., p. 14.
[421] Ibid., pp. 14-15.
[422] Ibid., p. 15.

strengths and weaknesses to the table. Why replicate human reasoning when you can transcend it? At the 1963 Michigan conference, one attendee, Dr. Kossack, reasoned:

> I always seem to get the role of being the devil's advocate, not that I think it is a bad idea to try to get the computer to follow the human process. When I talk with business people I say that I am worried that when we do these things that we will find that in the computer we will have programmed the coffee break because we've copied what people do. In fact, is there any evidence that the diagnostic process, as an art practiced by the medical profession, is better than chance?[423]

Kossack's comment here was in response to Lipkin's paper, where he suggested 'The results of our [the haematology group's] studies demonstrate the feasibility of automating certain portions of differential diagnosis'.[424] Lipkin answered to Kossack's provocation, observing that while it may be true that his haematology programme does not replicate the *exact* functions of the human mind (and that the ideal strategy may be to follow an alternative route) it is nevertheless worthwhile to 'give some thought' to the possibility that mind and machine, at least in the context of this particular program, 'bear some similarity'.[425]

At the Rockefeller Institute's 1959 conference on Diagnostic Data Processing, two of Lipkin's colleagues—Russell Ebald and Robert Lane similarly described the haematology 'matching' programme as a simulation of how doctors think. 'To be more specific', they declared, 'we attempted to simulate the thinking of the doctor'. Ebald and Lane proceeded to outline their view of how clinicians reason: 'Logically he would proceed as follows. Once the preliminary examination was complete and the results of appropriate tests had been recorded, the doctor would attempt to match these symptoms to some one disease or combination of diseases. He would do this either from his own

---

[423] Discussion of Lipkin, 'The Role of Data Processing in the Diagnostic Process', p. 276.
[424] Lipkin, 'The Role of Data Processing in the Diagnostic Process', p. 255
[425] Ibid.

experience or from the experience of his colleagues, or from texts of classical disease definition. This then is what we tried to do'.[426] That is, the team defined and computerised certain disease definitions and developed an algorithm to match the individual patient findings to these larger pictures—just as the physician 'matches' patient particularities to his own generalised definition of disease, created through his experience, study, colleagues, or some combination thereof.

In the discussion following Ebald and Lane's presentation, Engle pointed out other parallels between the programme and his notion of clinical cognition. As Engle saw it, physicians, 'when confronted with a large number of signs and symptoms, often picked from these one or a few key signs or symptoms, usually ones which were unquestionable. Then on the basis of this sign or symptom say, jaundice, he would automatically go into what we call a differential diagnosis, picking the diseases that would produce jaundice. Next the physician would consider these diseases in terms of all the signs and symptoms'.[427] Engle felt that the programme approximated this human process of thinking by, first, matching patient findings to a broad disease class on the basis of key symptoms and only later narrowing in on the more specific disease diagnosis.

Yet as the haematology project aged, the group came to side with Kossack's viewpoint. The haematology team presented this line of thinking in a 1975 paper:

> Like all other parts of the human body, the brain, the organ of thought, has become adapted to its environment through the evolutionary process. However, there are certain inherent limitations which restrict the development of organs and do not allow the optimum development of all aspects of function. For example, although animals have evolved legs and some have wings, none has evolved wheels, although we consider these essential for smooth and fast travel

[426] Ebald and Lane, 'Digital Computers and Medical Logic', p. 283.
[427] Ibid., p. 286

> on the earth. Man has developed tools of all kinds to improve and extend his strength and endurance.[428]

As with automobiles, the group suggested, so too with automated diagnosis. The computer offered ways not just of replicating human abilities but of building upon and extending those capabilities.

Of course not all developers of early computerised diagnostic systems agreed that the explication and replication of diagnostic reasoning was so hopeless—or fruitless. As is clear in Brodman's case, the process of formalising diagnosis forced physicians to think in new ways about the nature of the medical mind. Through his work computerising diagnosis, Brodman clarified his own ideas about clinical reasoning—ideas that he then codified into the software of his diagnostic programme. Others reported similar experiences. One crowning achievement of early efforts to computerise diagnosis, many later noted, was the explication and formalisation of the medical decision making process. The overlap between enterprises in engineering (building a diagnostic system) and those in psychology (clarifying the nature of clinical reasoning) were hard to miss. Many of the same people who undertook research initiatives in the computerisation of diagnosis also founded and populated some of the earliest medical societies and conferences dedicated to understanding the diagnostic process and to cultivating the emerging field of 'decision science'. Lee Lusted, for example, founded the Society for Medical Decision Making (SMDM), and conferences like those above on 'The Diagnostic Process' were teeming with the likes of Lipkin, Engle, Lusted, and Goldstein.

---

[428] Ralph L. Engle, Betty J. Flehinger, Scott Allen, Richard Friedman, Martin Lipkin, B. J. David, and Leo Leveridge, 'HEME: A Computer Aid to Diagnosis of Hematologic Disease', *Bulletin of the New York Academy of Sciences* 52, no. 5 (1976): 584-600, p. 585.

In his history of Logic and Heuristics Applied to Synthetic Analysis (LHASA), a computer programme designed to help organic chemists with retrosynthetic analysis, the historian Evan Hepler-Smith shows how the process of creating the programme 'shaped a way of thinking taken up by chemists, unaided by machines'.[429] That is, in the process of having to explicate their reasoning processes for the computer, the team defined new approaches to scientific reasoning that other chemists began to use outside the context of computerisation. Such observations parallel Suchman's observations (discussed in Part I) about the role of computers as 'powerful disclosing agents', as tools that physicians and others used to reflect upon, and better understand, themselves (or at least what they believed to be true about themselves).[430]

Similar phenomena were at play in the computerisation of diagnosis. The computer was an object that physicians began to think with, against, and alongside. The dynamics of 'powerful disclosure' were particularly decisive in the creation of INTERNIST-I, a programme at the University of Pittsburgh designed to make the majority of diagnoses in all of internal medicine. Developed under the auspices of the Stanford University Medical Experimental Computer for Artificial Intelligence in Medicine (SUMEX-AIM) (see Part III for more details), INTERNIST-I grew out of a collaboration in the early 1970s between the computer scientist Harry Pople and the internist Jack D. Myers. Myers had just stepped down from his perch as Chairman of the Department of Medicine and taken up the role of 'University Professor', a super-departmental appointment (modelled after similar professorships at Harvard) that gave faculty extraordinary levels of freedom to work and collaborate across disciplinary

[429] Evan Hepler-Smith, '"A Way of Thinking Backwards": Computing and Method in Synthetic Organic Chemistry', *Historical Studies in the Natural Sciences* 48, no. 3 (2018): 300-337, p. 300.
[430] Lucy Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions* (Cambridge: Cambridge University Press, 2006), pp. 226-240.

lines. Deciding where to direct his energies during this new phase of his career, Myers was drawn to a set of concerns and challenges that almost seemed to be nagging at the medical field's collective consciousness. First, how to manage a knowledge base that had grown 'vastly too large for any single person' to know, and second, how to ensure that the 'busy physician' considered 'the right answer particularly if the diagnosis was an unusual case'.[431]

Through a mutual colleague, Myers—who knew nothing about computers—was connected with Pople—who knew nothing about medicine. Together, they worked to develop a programme that would incorporate the medical expertise and imitate the logical processes of this single physician, Myers. Of the internists in the United States at the time, Myers had a strong claim to be The One to be computerised. Myers enjoyed an eidetic memory—and the reputation of being a 'master diagnostician'. Their collaboration thus began with Myers analysing—and performing—the diagnostic process before Pople. Pople's job was to watch, to formalise, and then to programme (Figure 20). Randolph Miller, an MD-PhD student who later joined the project, described the logic at play here: 'find a genius and can their brain'.[432] Myers described these sessions in greater detail: 'I chose a goodly number of actual cases from clinical pathological conferences (CPCs) because they contained ample clinical data and because the correct diagnoses were known. At each small step of the way through the diagnostic process I was required to explain what the clinical information meant in

[431] Jack D. Myers, 'The Background of INTERNIST I and QMR', *HMI '87 Proceedings of ACM Conference on History of Medical Informatics* (New York: ACM, 1987), 195-197, p. 195.
[432] Miller, quoted in 'Dr. Jack Myers, 84, a Pioneer in Computer-Aided Diagnosis', *New York Times*, 22 Feb. 1998.

context and my reasons for considering certain diagnoses. This provided to Pople insight into the diagnostic process'.[433]

Because of INTERNIST-I's reliance on Myers, Pople characterised the system as 'an example of simulation', its heuristic rules and logical processes simply those that best 'resembled the behavior [Pople] saw'. In being 'forced' to explain his diagnostic reasoning, Myers also felt himself changed. 'After analysing dozes of such cases', Myers wrote, 'I felt as though I had undergone a sort of "psychoanalysis"'.[434]

Even though Engle came to conclude that it was hopeless to have computers attempt to simulate the 'uniquely human ability' of 'intellectual synthesis', he and others in the haematology group acknowledged that their work with computers did much to clarify their own thinking about, and understanding of, the diagnostic process.[435] During the project's later twilight years, Engle insisted, perhaps departing from the perspectives of some of his earlier colleagues, 'the program was never intended to be a diagnosis machine'. Instead Engle defined the program's use in three other areas, one being its role in illuminating how physicians think: 'It would perhaps be better to ask whether or not there is a use for such a program. Three possible uses have been considered: 1) as a diagnostic aid, 2) as a teaching-learning device and 3) as a tool for studying the diagnostic process'.[436]

It was surely true that the haematology programme raised fundamental epistemological questions for many members of the group—particularly for Engle. Engle presented some of his reflections on these matters at the 1963 conference on

---

[433] Jack D. Myers, 'The Background of INTERNIST I and QMR', *HMI '87 Proceedings of ACM Conference on History of Medical Informatics* (New York: ACM, 1987), 195-197, p. 195.
[434] Ibid.
[435] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making', p. 216.
[436] 'HEME2', draft, 30 Mar. 1981, REP/11/8.

'The Diagnostic Process'. 'Work on the problem of computer aids to diagnosis', he noted, 'has forced us to try to define terms more precisely. It has also stimulated an interest in basic theoretical and philosophical questions which have some application to medical diagnosis; such questions as the theory of universals and the closely related problems of classification, definition, and nomenclature as well as induction and deduction'.[437]

Engle elaborated upon some of these points in a three-part series of articles published in *Archives of Internal Medicine* on the topic of 'Medical Diagnosis: Present, Past, and Future'.[438] The first of these articles, co-authored with B. J. Davis, dealt with the nature of diagnosis and the definitions of disease. Diagnosis, Engle and Davis argued, is essentially a two-sided coin: it refers both to the *process* of clinical decision making and to the *decision* itself. Its functions, moreover, are multiple, involving not just 'the naming and classification of disease' but also the communication of implicit theoretical frameworks about health and illness and 'the construction of hypothesis by serving to stimulate investigation'.[439] Engle and Davis proceeded to show how disease definitions can be just as amorphous as the concept of diagnosis. 'At the present time', they observed, 'there is no unified concept of disease and, therefore, no unified yet flexible and practical concept of diagnosis'.[440] Instead, they argued, the nature of disease concepts varies considerably—ranging from diseases whose aetiologies are

---

[437] Ralph L. Engle, *The Diagnostic Process*, p. 12.
[438] Ralph L. Engle and B. J. Davis, 'Medical Diagnosis: Present, Past, and Future. I. Present Concepts of the Meaning and Limitations of Medical Diagnosis', *Archives of Internal Medicine* 112 (1963): 512-519; Ralph L. Engle, 'Medical Diagnosis: Present, Past, Future. II. Philosophical Foundations and Historical Development of Our Concepts of Health, Disease, and Diagnosis', *Archives of Internal Medicine* 112 (1963): 520-529; Engle, 'III. Diagnosis in the Future, Including a Critique on the Use of Electronic Computers as Diagnostic Aids to the Physician'.
[439] Engle and Davis, 'I. Present Concepts of the Meaning and Limitations of Medical Diagnosis', p. 516.
[440] Ibid., p. 516.

well understood and whose clinical presentations are uniform (such as certain congenital cardiac malformations), to diseases whose aetiologies are not at all known and whose clinical presentations are highly idiosyncratic (such as 'collagen diseases', 'aplastic anaemia', and other more syndromal disease categories).

The second article in the series turned largely to matters of philosophy. Engle argued that the process of diagnosis represented an instantiation of the philosophical concept of 'the theory of universals', with physicians having to constantly negotiate and navigate between the particular (the individual patient) and the general (the disease concept). 'Every time a physician thinks about a patient or makes a diagnosis, he must consider the patient as an individual and at the same time generalize as much as possible about his condition so that it can be classified and a diagnosis made'.[441] Engle called the back-and-forth required of the physician a 'psychic shuttle', borrowing a term from the English critic Herbert Read.

In the final instalment of the series, Engle, finally, addressed the computer's role in all of this. Engle saw a clear role for the computer in medicine. It was a tool that boasted a unique suite of strengths: 'We may conclude that electronic computers, because of their large data-storage capacity or memory, their accuracy and speed of data recall, and their ability to process rapidly large volumes of information, should be of great value to physicians for aiding in certain aspects of the thought processes required for medical diagnosis and even in aiding decision concerning therapy'.[442] Yet he was careful to note that he saw its role as a limited one. Engle's reflections on the diagnostic process had convinced him that diagnosis included both deductive

[441] Engle, '. II. Philosophical Foundations and Historical Development of Our Concepts of Health, Disease, and Diagnosis', p. 521.
[442] Engle, 'III. Diagnosis in the Future, Including a Critique on the Use of Electronic Computers as Diagnostic Aids to the Physician'.

(reasoning down from general principles) and inductive (reasoning up from specific observations) forms of reasoning, and that the inductive modes of reasoning in particular contained an 'intuitive' component that defied computerisation: 'Insofar as our inductive thinking is based on statistical and probability principles, the computers will be of great help in compiling essential data. The role of intuitive induction in diagnosis has not been clarified. It is my opinion that intuitive induction cannot be performed by a machine and will remain solely within the physician's province'.[443]

BACK TO BAYES

The haematology group could not escape the fatal magnetism of Bayes' theorem. By the late 1960s, the group had abandoned their 'matching' approach, citing the method's 'severe limitations'.[444] In its place they returned to a method they had once eschewed: they developed a programme called 'HEME', which operated by a modified Bayesian algorithm.

This shift in strategy in many ways reflected the group's shifting leadership. While the project's nucleus stayed in New York City throughout its existence, it bounced among different institutional homes. After Zworykin relinquished leadership in 1961, the project moved from the Rockefeller Institute to the New York University College of Engineering, where it came under the direction of the statistician Max Woodbury (1961-1966) and later the sociologist Louis J. Gerstman (1966-1967). In 1967, the project moved from the NYU campus in Midtown Manhattan to Brodman's stomping grounds in the Upper East Side. There, at Cornell University Medical

[443] Ibid., p. 541.
[444] 'Critique on Computer Aids to Medical Diagnosis', n.d., REP/4/3.

College, the project fell under the supervision of Ralph Engle who was joined at the helm by the mathematician Dr. Betty Flehinger a year later.

Flehinger was *the* driving force behind the model's substantive changes. Trained as a biostatistician, Flehinger worked as a research assistant in Columbia's Mathematics Department before joining the university's IBM Watson Laboratory in 1959. She maintained this affiliation throughout her work on the haematology programme. Flehinger's reputation preceded her. On learning of Flehinger's appointment, Woodbury wrote to Engle, 'I am glad to see you have the cooperation of Betty Flehinger in your work. She is very good from what I know about her'.[445] Flehinger delivered on these expectations and quickly began steering the group's research direction. The group spelled out her contributions in a 1972 grant application: 'Because of our interest in the "significance model" we did not start working with the Bayesian model until Dr. Betty Flehinger joined our group in 1967'.[446] Making the transition to the new model even more palatable, Flehinger noted that the same data that had been amassed for the significant model could be used to estimate the probabilities required for Bayes' theorem. 'Dr. Flehinger', the team wrote, 'pointed out that using most of the information accumulated for the significance model we would be able to apply a new version of Bayes' theorem, one which does not assume mutual exclusiveness of diseases'.[447]

Classically, a Bayesian model operates under the assumption that the patient in question only has one disease. The absurdity of this assumption drove many away from the theorem. But Flehinger proposed a work-around. Whereas classical Bayesian

[445] Woodbury to Engle, 12 Mar. 1970, REP/4/2.
[446] Grant Application, Department of Health, Education, and Welfare, 'Trial of a New Bayesian Decision System for Diagnosis', 1 Oct. 1972 (HL 15927-01), REP/5/3.
[447] Ibid.

models took multiple diseases into consideration at once—assuming a universe of patients where each patient has but a single disease—the modified model only considered one disease at a time. The assumed universe of the modified algorithm also looked different, consisting of two groups of patients: those with the disease in question and those without the disease in question (whether healthy or suffering from other diseases). Under this new model, the estimated probabilities of all diseases need not add up to one; that is, the model allowed for coexisting diseases in a patient. HEME, operating under this model, could iteratively compute the probability that the patient has each of the 40 diseases presently considered by the system. The system could also marshal for the physician-user the findings that supported, and those that opposed, each of the diagnoses on the differential.

Another selling point for the new model was what Engle and Flehinger called its 'self-improving' nature.[448] As the computer saw more and more patients, it could incorporate this data and update its underlying conditional probabilities. With more and better data, the system should improve over time with experience—not unlike a practicing physician. The model wasn't just self-improving though; the team continued to add their own improvements as well. For one, the team continued to update the programme's database, which expanded over time to encompass 142 hematologic diseases and 667 findings.

The team also modified fundamental aspects of how the programme worked. As early as 1974, the team had already started creating a new version of the program, HEME2. These changes were responsive to the team's investigations into the original HEME programme's deficiencies: 'After analysing a number of diagnostic situations in

[448] Betty J. Flehinger and Ralph L. Engle, 'HEME: A Self-Improving Computer Program for Diagnosis-Oriented Analysis of Hematologic Diseases', *IBM Journal of Research and Development* 19, no. 6 (1975): 557-564.

which the computer went astray in the diagnostic process we decided that part of the problem related to the fact that HEME was a one-step system'.[449] HEME2 functioned by two-steps. In the first step, the programme classified the patient into one or more of nineteen broad categories of haematologic disease; in the second, the programme diagnosed the patient as having one of the diseases within the identified categories.[450] By classifying patients, first, into broad categories of disease and, only later, into the actual diagnosis, HEME2 in many ways carried forward the new ways of considering and classifying haematologic diseases that had been necessitated earlier on by the machine's material affordances.

Going back to Bayes', though a well-trodden path, presented its own challenges. Most pressingly, the approach remained vulnerable to many of the same critiques that members of the haematology team had levelled twenty years earlier. Woodbury, who had articulated the group's early critiques of Bayesian approaches, found the new, modified formulas developed by Flehinger nearly as restrictive as earlier ones applied by Ledley and Lusted. Woodbury wrote in a note to Lipkin, describing his preliminary thoughts from his 'investigation of the conditions for validity for the use of Flehinger's formulas', '[t]he conditions are quite severe and the more general formulas required to remove them are considerably harder to use …'.[451] He expounded upon his concerns in a letter to Engle. 'I am a bit puzzled by the notes

[449] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making', p. 214.
[450] As an example, 'aplastic anemia' was one of the nineteen 'categories', under which there were the following 'diseases': 1. Fanconi's Anemia, 2. Acquired Idiopathic, 3. Secondary Drugs, Chemicals, or Irradiation, 4. Secondary to Infection or Immunologic Abnormalities, 5. Myelophthisic Anemia, 6. Paraoxysmal Nocturnal Hemoglobinuria, 7. Congenital – Without Deformities (Estren and Dameshek). Engle and Flehinger, 'HEME2: A Long History of Computer-Aided Hematologic Diagnosis', in Bruce I. Blum (ed.), *Proceedings of the Six Annual Symposium on Computer Applications in Medical Care* (Computer Society Press, 1982), 763-767, p. 765.
[451] Woodbury to Lipkin, n.d., REP/4/2.

on the Bayesian Model used. It seems that you must have assumed that each finding was unique to a disease with what is, essentially, a uniform probability of "false" possible for the finding for all other diseases. If this is even approximately true it represents an enormous simplification of the problem'. Woodbury had other concerns, too, but he merely directed Engle to his earlier paper on the 'Inapplicabilities of Bayes' Theorem to Diagnosis': 'Other than this', he wrote, 'my criticisms are all pretty minor and are inferable from my article on "Inapplicabilities…"'.[452]

Reviving another line of criticism against Bayes' theorem, others found HEME's new strategy questionable on the basis of its data requirements. In 1974, Engle and Flehinger lost a grant bid, in part because of these unresolved issues. Channelling many of the concerns about data judgement and bias discussed earlier, the grant reviewers felt that the group had not paid adequate attention to matters of data quality. 'The applicant', they explained, 'did not demonstrate that sufficient thought had been given to the critical issue of validating the input data'.[453]

Even if the quality of input data could be assured, the reviewers were sceptical that the group could harvest the appropriate quantity. The back-of-the-envelope mathematics didn't look promising. They broke it down: 'A major reservation concerned the practicality of generating useful conditional probabilities for a system containing 120 diseases. With over 500 findings at present and perhaps 1,000 findings in the expanded system, the application would need to generate approximately 100,000 conditional probabilities. The number of cases required to estimate these conditional probabilities could number in the millions'. Alternatives to collecting such masses of 'raw' data—namely, 'textbook review or physician panel'—would still constitute a

[452] Woodbury to Engle, 12 Mar, 1970, REP/4/2.
[453] Dorothy Reese to Engle, 2 Aug. 1974, REP/5/4.

'tremendous effort' of their own. (One expects Atamer would have agreed.) The outcome of all this work, moreover, 'would only apply to patients known to have hematologic disease'—a small fraction of patient diagnoses.

A reviewer for the *American Journal of Medicine* (*AJM*) criticised the HEME project in 1975 on similar grounds, writing, 'One of the major limitations of this diagnostic approach has been the lack of availability of statistical data'. Of the data the researchers could get, the reviewer questioned their quality—and origin: 'The title advertises a "self-improving" program. This is no[t] true—there is only a proposed scheme for updating the data base manually substituting better data. This method depends on the accuracy to which ni and su are known, both of which are "pulled out of air"'.[454]

Apart from these well-known challenges, a Bayesian approach to computerising diagnosis had, by the late 1960s and early 1970s, lost the sharp edge that it once enjoyed. As new faces in the field pioneered 'expert systems' and self-avowedly 'AI' approaches to diagnosis (see Part III), the Bayesian method came across as a bit outdated, ho-hum, even regressive. This view set the tone of the 1975 *AJM* referee report: 'The manuscript is a rather typical example of the Bayesian statistical diagnosis programs that have been appearing for the past fifteen years. Although it purports to be different than prior attempts, it is not. It falls into the same old traps and limitations'.[455] The reviewer thus concluded, 'In summary, this manuscript would have been interesting 5-10 years ago, but now is just one of many like it. Since it addresses no new issues and since the program remains untested I question its interest to *AJM* readers'.[456]

---

[454] William Grace to Flehinger, 29 Jan. 1975, REP/4/2.
[455] Ibid.
[456] Ibid.

These troubles above the water reflected other troubles below. In their publications, the team boasted of HEME's 'encourag[ing]' results during testing.[457] This outward optimism, however, belied ongoing uncertainties and challenges. The team had the frustrating sense that they had reached a point of diminishing returns: edits to the programme solved some issues, only to introduce others. Modifications to the programme and its database therefore frequently resulted in no net performance gains. 'In HEME2 our attempts to correct the deficiencies were directed toward modifying the subjective information in the data base of the program', they explained. 'In doing so, we found that these changes resulted in improving some parts of the system, often at the expense of deterioration in other parts'.[458]

The asymptote of performance that the HEME programme seemed to have reached wasn't even that impressive. The system did fine when fed simple cases. But when presented with more complex cases—the precise kind of cases for which physicians might need computerised decision support—the system seemed prone to choke. The group noted in a clear-eyed assessment of HEME: 'If given relatively simple cases, the program does very well. If given very difficult cases it does poorly and it is sometimes difficult to be certain what the correct answers should be. Usually, however, the program gives a reasonable, though not necessarily correct answer. In moderately difficult cases, its performance is variable.'[459] Thinking of these challenges, Engle thus concluded, '[I] don't believe we have evidence yet that "accuracy often

---

[457] Tested on a series of 31 cases from the medical records of New York Hospital, HEME was found to have diagnosed 14 cases as 'excellent', 7 as 'good', 1 as 'fair', 3 as 'poor', and 6 as 'not evaluable' by a panel of expert haematologists. ('Not evaluable' simply indicated that HEME did not include the ultimately diagnosis in its repository.) Engle et al., 'HEME: A Computer Aid to Diagnosis of Hematologic Disease', p. 595.
[458] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making', p. 215.
[459] 'HEME2', draft, 30 Mar. 1981, p. 14, REP/11/8.

approaches that of a skilled hematologists." We believe it is premature for this system to have any impact on clinical hematology'.

But Engle and Flehinger wondered: if not clinical haematology, perhaps there was another suitable home for the programme. As they retreated from the hope that HEME would find its way into the rhythms of clinical care, they increasingly worked to define a new niche for the programme in medical education. Toward this end, they introduced HEME as a teaching tool into the haematology elective for fourth-year medical students at Cornell University Medical College. The idea was that by working through a patient case alongside HEME—comparing one's own reasoning with that of the computer—early learners would come to better understand the intricacies of haematologic decision making and the fundamentals of the diagnostic process.

Yet even medical students—a group happy to mindlessly grind through curricular unpleasantries—resisted the programme. Of the eight students invited to use the programme, only four ended up using it and only one finished the syllabus.[460] The reasons behind this degree of attrition were variable and many, ranging from '[o]bvious scepticism by many on the senior staff concerning the ability of a computer to aid the teaching' and delayed responses times due to 'computer difficulties', to students being 'very busy with their requirement work' and 'unfamiliarity in the use of computers'.[461] The matter ultimately boiled down to a matter of cost-benefit analysis for the students: 'the students found it to be too unwieldy and time-consuming to sustain their interest and effort in using it'.[462]

---

[460] Engle and Flehinger, 'Report on Work Conducted Under the Biomedical Research Support Grant', 22 Jun. 1977, REP/6/12.
[461] Ibid.
[462] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making', p. 214.

The team had high hopes for HEME: 'It is hoped that the program ultimately will gain acceptance by physicians and will be used routinely as a step in the diagnostic process'.[463] Yet by the 1980s, it was clear that this project, now ongoing for nearly three decades, would not gain clinical traction.[464] HEME had failed: it failed in its earlier quest to become a widespread aid in differential diagnosis, and it even failed in its later reformulations as a teaching aid. 'For twenty-five years', the team reflected, 'investigators have labored enthusiastically toward the development of computer aids to diagnosis. … At this point we have decided to end our labors in this field…'. Having dedicated such a sizeable chunk of his career to the project, Engle was introspective about the possible reasons behind these persistent challenges. The haematology group's failures were surely overdetermined. Among the factors contributing to the project's demise, Engle identified poor user-interfaces, unruly computing technologies, limitations of the program, and professional resistance to computerised diagnostic aids. Yet Engle concluded such factors, while contributory, were not really 'critical'.[465] Other factors, he believed, were steering the ship.

Some wondered, for example, whether computerised systems like HEME were really geared to the right problem—whether these systems excessively fetishised the act of 'diagnosis' (assigning labels) at the expense of 'decision making', which they saw as the broader driving force of patient care. Dr. William Bauman articulated this view in a 1969 letter to the editor of the *New England Journal of Medicine*: '"Diagnosis" is the act of identifying a disease or giving a label to a set of medical findings. The major

[463] Engle et al., 'HEME: A Computer Aid to Diagnosis of Hematologic Disease', p. 598.
[464] The group disbanded in 1972, but Flehinger and Engle continued their collaboration into the late 1980s.
[465] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making', p. 216.

objective in patient care, however, is not "Diagnosis" but rather "Management." A patient does not benefit from the label; he is helped by the physicians' recommendations, or actions. Nor is the diagnosis necessary to make a decision as to management. In fact, it is entirely possible to [deliver] good medical care without ever identifying the name of the illness'.[466] Computerised systems, Bauman implied, should be less fixated on diagnosis. Instead they should take up the broader remit of helping physicians make better decisions of all kinds as they manage their patients. By this reasoning, such efforts may have failed because they weren't appropriately geared toward solving the more pressing problems in medicine. 'Our major concern', Bauman concluded, 'should be how to best treat the patient, not how to make a diagnosis'.[467]

As Engle reflected on the fate of his work, he endorsed the possibility that programmes like his were just not filling a real need. No matter how good the programme got, perhaps it was not sufficiently well aligned with the needs of its users. Enumerating some of the reasons HEME 'has received little use', Engle and Flehinger wrote: 'It is rare that a purely diagnostic problem in which a name is to be given to a patient's ailment presents itself. Most of the physician's time is spent in the day-to-day management of patients of known diagnosis'.[468] Even on the rare occasion when a narrow diagnostic task presents itself, it often involves practices and tools that defy computerisation. They noted, 'Many of the most important diagnostic tools, such as

---

[466] William Bauman to Franz Ingelfinger, 11 Jul. 1969, REP/8/2.
[467] Ibid. Bauman was in fact enthusiastic about the possibility of applying computers to patient management. He simply saw the singular fixation on 'diagnosis' as an iceberg; he was keen to redirect the ship. He wrote to Engle, 'I believe that substantial proportion of patients' problems can be managed by the computer…better than they are now managed by the friendly druggist, the articulate TV announcer, the school or screening nurse, and even by some of our hard-working colleagues. I am not altogether sure that a disease diagnosis is needed. The cluster of findings, unlabeled, may be all that is required to give sound management advice'. Bauman to Engle, 1 Dec. 1969, REP/8/2.
[468] Engle and Flehinger, 'HEME2: A Long History', pp. 764-765.

palpation of the spleen, bone marrow and blood smear examination, and general observations of the patient cannot be transferred to a computer. They require unique human pattern-recognition abilities'.[469]

Ultimately, though, Engle's thirty years of experience working on the programme led him to the conclusion that the greatest obstacle to computerised diagnostic systems was nothing external to, or modifiable within, medicine. It was the nature of medicine itself—and of the diagnostic mind. Engle explained,

> In our estimation, the critical impediment to the development of decision programs useful in medicine lies in the impossibility of developing an adequate data base and an effective set of decision rules. Any decision module contains only a small fraction of the data and of the intricate relationships that a physician learns in the course of education and practical experience. Personal experience over many years trying to develop a suitable diagnostic program for hematologic diseases confirms the great difficulty and even the impossibility of incorporating the complexity of human thought into a system that can be handled by a computer. A computer-based diagnostic system has three major components, each of which defies successful implementation: input, inference by means of an algorithm, and output.[470]

At every level, Engle believed, medicine and the mind were too slippery, too ill-defined, too complex.

* * *

In his three-part series on 'Medical Diagnosis', Engle compared disease to light. By the Copenhagen interpretation, light can behave either like a wave or like a particle—depending on the questions one is asking. Similarly, disease definitions can take on different properties and characteristics—depending on one's frame of mind. Engle explained,

> At its heart, the question of qualitative disease versus quantitative disease is not unlike the quantum theory versus the wave theory of light. Either hypothesis can be invoked if need arises, and in this sense both are correct. At the most

[469] Ibid., p. 765.
[470] Engle, 'Attempts to Use Computers as Diagnostic Aids in Medical Decision Making', p. 216.

> fundamental level in the chromosome or the DNA, we now think of changes as being discontinuous or qualitative. Nevertheless, because of the long chain of events between the chromosomes and the manifestations that we can interpret, the many interactions that occur at these stages produce effects which are quantitative.

There are different ways of defining and conceptualising disease entities. A purpose of this section has been to demonstrate the role of computers—and material culture more generally—in shaping how we see and define disease entities.

This section, read together with the preceding section, thus yields insights into the process of disease definition—and the participation of computers and material culture in that process. Sometime in the 1960s, Ralph Stacy, an expert in biomedical computing at Ohio State University, proclaimed that the computer's future in medicine looked bright. '[C]omputers', he forecasted, 'will play a very real part in the medicine of the future'.[471] The basis of Stacy's conjecture was the increasing quantification of medical knowledge: 'The nature of medical science is undergoing a change such that it is becoming more and more quantitative'. The stories of the MDS and the haematology programme together point to a more complicated interplay between computers and quantification than the kind suggested by Stacy. Where there are numbers, Stacy reasoned, computers may go. That is certainly true, but so is this: Where there are computers, numerical concepts may emerge. The arrows of causality point in both directions.

These cases thus illuminate the mechanisms by which computers have encouraged numerical diagnosis and numerical concepts of disease. As two of the earliest efforts to computerise medical diagnosis, they reveal that at the moment computers entered the arena of medical diagnosis, researchers began to collate

[471] Vladimir Zworykin, 'The Need for Electronics in Medicine', n.d., VZP/82/16.

knowledge and define disease in new ways. As soon as Zworykin and his colleagues started work on a diagnostic programme in haematology, they began to generate quantitative descriptions of disease. Similarly, in the process of creating the MDS, Brodman translated subjective symptoms into the language of numbers and thresholds. The diagnosis of disease, in both cases, was the result of an algorithmically guided comparison between a numerical rendering of the patient's health status and population health data that was stored in the computer's memory. Together the cases highlight the co-production of computers in medicine and new conceptions of disease. The development of computerised diagnostic systems resulted in the reshaping of disease categories. In a reciprocal fashion, the introduction of certain material methods of collecting patient data also led to the creation of computerised diagnostic systems.

PART III

Black Box Medicine

- CHAPTER V -

MYCIN Explains Itself

Over four days in the summer of 1975, pioneers in the computerisation of medical decision making gathered in New Brunswick, New Jersey to attend the First Annual Artificial Intelligence (AI) in Medicine Workshop, hosted at Rutgers University.[472] Present at the event were nearly thirty physicians, computer scientists, and mathematicians, affiliated with a range of different early computerised decision support projects. Among them, Edward Shortliffe, Stanton Axline, and Randy Davis attended as the frontmen of MYCIN, a consultation system designed to aid physicians in the identification and treatment of bacterial infections; Harry Pople, Jack Meyers, and Randolph Miller brought to the workshop their experience developing the expert system called INTERNIST-I, an ambitious system intended to recognise the majority of possible diagnoses in internal medicine; and Saul Amarel, Casimir Kulikowski, and Aran Safir discussed their work designing CASNET (causal-associational network), a programme designed to assist in diagnosing optic nerve diseases. These systems represented some of the first attempts to apply self-avowedly AI techniques to medical decision making. They were, some at the time believed, 'the only prototypes of biomedical knowledge-based systems that had demonstrated the feasibility of exploiting large quantities of structured domain knowledge to solve complex problems in clinical consultation'.[473]

---

[472] Saul Amarel, Casimir A. Kulikowski, N. S. Sridharan, and Deirdre Sridharan, *Proceedings of the First Annual AIM Workshop, June 14-17, 1975*, JLP/94/25.
[473] Casimir A. Kulikowski, 'Knolwedge-Based Systems in Biomedicine: A 10 Year Retrospective', *Proceedings of the Annual Symposium on Computer Application in Medical Care* (1986): 423-424, p. 423.

Beyond representing some of the most promising—and hyped—efforts of their time, the systems present at the workshop were unified by a more immediate and tangible connection: they were developed under the auspices of the Stanford University Medical EXperimental Computer for Artificial Intelligence in Medicine (SUMEX-AIM). Beginning in January 1974, SUMEX-AIM was established as a national computing facility at Stanford to support research on medical applications of artificial intelligence. The facility's nature and scope was outlined in a June 1974 prospectus: 'The SUMEX … facility, under Professor Joshua Lederberg as Principal Investigator, is national in scope in that a major part of its computing capacity will be made available to authorized research groups throughout the country by means of a communications network'.[474] This 'communications network' was none other than ARPANET—the Department of Defense-sponsored forerunner to the Internet. Indeed, SUMEX-AIM was the first non-Department of Defense host on DARPANET's 63-node network.[475] In addition to describing SUMEX-AIM's national scope, the prospectus laid out its aims: 'There are two main objectives of the facility: 1) the specific encouragement of applications of artificial intelligence in medicine (AIM) and 2) the managerial, administrative, and technical demonstration of a national shared technological resource for health research'.[476] The Rutgers workshop, then, was an occasion to celebrate SUMEX-AIM's early strides.

---

[474] 'SUMEX-AIM Prospectus', 22 Jun 1974, p. 1, JLP/62/49.
[475] For discussion of SUMEX-AIM's configuration with ARPA, see, e.g., Tom Rindleisch to 'SUMEX File', 'Notes on First AIM Executive Committee Meeting (11/26/73), 3 Dec. 1973, JLP/62/34. See also: Edward H. Shortliffe, 'The Changing Nature of Telecommunications and the Information Infrastructure for Health Care', in *The Changing Nature of Telecommunications/Information Infrastructure* (Washington, D.C.: National Academy Press, 1995), 67-73, p. 67.
[476] 'SUMEX-AIM Prospectus', p. 1.

It was also an occasion to discuss common topics, challenges, and themes that cut across SUMEX-AIM-supported projects. The workshop broached topics that covered a spectrum that ranged from the technical to the philosophical. One topic in particular seemed to span both ends of this spectrum: the issue of communicability. How could advice generated by computerised algorithm be communicated—and made understandable—to human medical personnel? To many, this question begged a technical answer: communicability was a problem to be solved through technological innovations in natural language processing and human-computer interfaces. But philosophical considerations drove discussions around interpretability and communicability as well. To what extent were communicability and interpretability driving values in medicine and medical decision making? Could they justifiably be sacrificed at the altar of other medical values—say, those of speed and accuracy?

Using the MYCIN system as its primary case study, this section explores the interrelations among computing, authority, and trust in post-war American medicine. It demonstrates that concerns about trust and transparency drove much of the development of, and responses to, the MYCIN system. Moreover, concerns about the extent to which a computer's advice could be understood, communicated, and trusted co-existed with concerns about authority and the physician's identity. This section thus elucidates a dynamic that structured the creation and reception of computerised diagnostic systems through the second half of the twentieth century—a dynamic that endures today. This story begins in the arid foothills of the South San Francisco Bay Area, at a university that was just coming into its own.

COMPUTING AT THE RANCH

Today no university is associated more strongly with the fields of medical informatics, biomedical computing, and computer science than Stanford University (Figure 21). Stanford's dominance in these areas, however, was not foreordained; indeed, in the 1940s and 1950s the university looked nothing like an optimal incubator for innovative work at the intersections of biology, medicine, and computing.[477] Plagued by an inadequate flow of research funds, a medical school that rested on financially shaky grounds, and a research culture characterised by isolation and siloing, Stanford in the immediate post-war period seemed to lack the basic ingredients required for a costly research enterprise that sought to bridge the computational and the medical sciences. But by the 1960s, all of this had changed; against the odds, Stanford had emerged as the epitome of an integrated and interdisciplinary research university with a healthy share of federal funding.

This transformation was less the product of some natural or serendipitous evolution than a conscious effort to restructure and revitalise the university. Such an undertaking was spearheaded by the engineer and administrator Frederick Terman, who served as Stanford's Academic Dean beginning in 1946 and as its Provost from 1955 to 1965.[478] From these perches, Terman laid the groundwork for the institution's transformation into a STEM heavyweight. He attracted major grants from the Department of Defense; he expanded the university's departments in the fields of

---

[477] November, *Biomedical Computing*, p. 222. This section's discussion of Stanford, ACME, and DENDRAL draws on November's excellent work.

[478] On Frederick Terman and Stanford's post-war transformation, see, e.g., Stuart W. Leslie, 'Playing the Education Game to Win: The Military and Interdisciplinary Research at Stanford', *Historical Studies in the Physical and Biological Sciences* 18, no. 1 (1987): 55-88; Rebecca S. Lowen, '"Exploiting a Wonderful Opportunity": The Patronage of Scientific Research at Stanford University, 1937-1965', *Minerva* 30, no. 3 (1992): 391-421; Rebecca S. Lowen, *Creating the Cold War University: The Transformation of Stanford* (Berkeley, CA: University of California Press, 1997); Cyrus C.M. Mody, *The Long Arm of Moore's Law: Microelectronics and American Science* (Cambridge, MA: MIT Press, 2017), pp. 25-27.

science, engineering, and statistics; and he forged new linkages between the university and the tech industry. Especially noteworthy among Terman's initiatives was his founding of the Stanford Industrial Park (later Stanford Research Park), an arrangement whereby Stanford leased its land to a number of technology companies like Hewlett-Packard, Lockheed Corporation, General Electric, and the Shockley Transistor Laboratory. By catalysing these linkages, Terman helped set into motion a series of developments, spanning the industry-academy divide, that would culminate in the rebranding of 'Santa Clara Valley' into 'Silicon Valley'.[479]

Important changes particular to the School of Medicine and the Department of Computer Science paralleled this larger institutional transformation. The School of Medicine, once financially strained and physically isolated, had, by 1959, found a firmer financial footing and moved from its former site in San Francisco to the main campus in Palo Alto. A key motivation behind this relocation to the university's main campus was the hope that the School of Medicine's physical proximity to Stanford's larger research operation in the biological, physical, and engineering sciences would promote cross-fertilisation between fields. As explained in a 1965 grant application for research to be conducted at the School of Medicine, 'Stanford made a calculated decision to move the medical school from San Francisco to the Palo Alto campus, and did so in 1959. This was a conscientious self-dedication to far-reaching communion of medical research and education with university life, based on the principle of mutual interdependence of medicine with the physical sciences and with human affairs'.[480]

---

[479] The name 'Silicon Valley' was popularised by in series of articles, titled 'Silicon Valley in the USA', for the weekly trade newspaper *Electronic News.* Don C. Hoefler, 'Silicon Valley in the USA', *Electronic News*, 11 Jan. 1971.

[480] Joshua Lederberg, 'Advanced Computer for Medical Research', Application for Research Grant, no. FR 00311-01, 4 Oct. 1965, Department of Health, Education and Welfare, Public Health Service, JLP/60/25.

Atop the list of departments with which researchers at the School of Medicine would form new and significant collaborations was the Department of Computer Science, founded in 1965.[481] The Department's rise to international prominence was short and swift. Its founding chair, George Forsythe, aggressively (and successfully) recruited some of the best and the brightest minds in computer science to Palo Alto. Under Forsythe's leadership, the Department attracted an enviable group of talented researchers, from established leaders like John McCarthy to promising up-and-coming scholars like Edward Feigenbaum.

It was within this atmosphere that the university made its earliest forays into the area of biomedical computing. The most noteworthy of these early efforts was the creation of the Advanced Computer for Medical Research (ACME), a computer system 'designed to provide powerful computing to research laboratories in the Medical School'.[482] The vision behind the AMCE was in many ways set by Forsythe, who saw computer science not as a mere sub-field of mathematics but as its own discipline with manifold lessons and applications for manifold other fields of inquiry. In pursuit of this broader vision, Forsythe worked vigorously as department chair to apply computers to biomedical research at Stanford, not only hiring international leaders in biomedical computing like Gio Wiederhold but also exerting great effort through the early 1960s to secure NIH funds for a large computer to support research in biology and medicine. As it turned out, Forsythe's personal attempts to acquire such a computer never

[481] Before the Department of Computer Science developed into its own department, it existed as the Division of Computer Science within the Mathematics Department. George Forsythe founded the Division of Computer Science at Stanford in 1961. November, *Biomedical Computing*, 228.
[482] Gio Wiederhold, 'A Summary of the ACME System', Presentation Given at the ONR Computer and Psychobiology Conference, the U.S. Navy Postgraduate School, Monterey, CA, 17 May 1966, JLP/60/45.

materialised. They did, however, set the stage for an even larger push to bring computers to bear on biomedical research that ultimately would yield positive results.

The ACME was, like many pioneering mid-twentieth-century enterprises in science and medicine, the brainchild of Joshua Lederberg, a Nobel laureate famous for his discovery of bacterial conjugation (Figure 22). Lederberg's path to the field of biomedical computing was circuitous: a geneticist, he arrived to the fields of biomedical computing and artificial intelligence by way of an intermediary fascination with extra-terrestrial life. On 4 October 1957, Americans across the country awoke to jarring news: the Soviet Union had successfully launched into orbit the world's first artificial satellite, besting the U.S. space program. As the United States, with its Cold War Manichean habit of mind, ramped up its space research programmes in response, Lederberg, then thirty-two years old and already a world-famous scientist, 'set out to assure that fundamental biological science was properly represented' in these programmes.[483] Initially Lederberg's capacity in this regard was purely advisory. But within a few years, Lederberg's role had shifted from that of advisor to that of researcher; in 1960, NASA established the Life Sciences Research Office and tapped Lederberg to create and head the Instrumentation Research Laboratory (IRL) at Stanford. Tasked with helping design the technical and conceptual foundations of space research, the IRL carried a rather broad and ill-defined research mandate. But Lederberg's interests soon stabilised around questions of extra-terrestrial life.[484] Was there life on Mars? And if so, how might it be located, isolated, and analysed? Might

---

[483] Lederberg, 'How DENDRAL Was Conceived and Born', ACM Symposium on the History of Medical Informatics, National Library of Medicine, 5 Nov. 1987, pp. 3-4, JLP/86/77.

[484] On Lederberg and contemporaneous research on extra-terrestrial life, see Audra J. Wolfe, 'Germs in Space: Joshua Lederberg, Exobiology, and the Public Imagination, 1958-1964', *Isis* 93, no. 2 (2002): 183-205.

Martian life present risks of disease and contamination to life on Earth? Pursuit of such questions brought Lederberg into confrontation with the possibilities of applying computers to research in the life sciences. As Lederberg would later relate, 'Both the internal activities of the Instrumentation Research Laboratory, and design discussions with the engineering managers of spaceflight missions (principally at Caltech's JPL) brought us into intimate conversation with technology of automation, process control, communications and computer management'.[485]

Having whet his palate, Lederberg 'quickly succumbed to the hacker syndrome'.[486] He became captivated by the technical aspects of computer programming in addition to its broader promise for the biomedical sciences. Lederberg's enthusiasm did not go unnoticed. Early in 1965, Terman and the School of Medicine's Dean formed the university's Computer Policy Committee (CPC) and appointed Lederberg as its chair. In this administrative role, Lederberg worked with tireless zeal to set up the ACME, define its parameters, and shore up its financial, technical, personnel, and conceptual foundation. Within months, by the fall of 1965, he was finalising a proposal to the NIH for the computer and 'pushing hard to get ACME working by [the following] May'.[487] Lederberg and his colleagues felt that the ACME would resolve an important unmet need. It did not, however, enter totally uncharted territory. Prior to the advent of the ACME, Stanford boasted established channels of computational support. Through the Stanford Computation Center (SCC), Stanford researchers had access to full-time service on two computers: an IBM 7090/1401 and a Burroughs 5500.[488]

---

[485] Lederberg, 'How DENDRAL Was Conceived and Born', p. 4.
[486] Ibid., p. 4.
[487] Lederberg to John Z. Bowers, 21 Sep. 1965, JLP/60/43.
[488] The IBM 7090/1401 supported the programming languages SUBALGOL, LISP and FORTRAN; the Burroughs 5500 supported ALGOL.

Yet even with the SCC, the ACME team believed that there remained a sizeable gap. Some of the SCC's shortcomings were general, the consequence of stretched human and technical resources that were unable to offer complete problem-solving services or adequate individual support. But other shortcomings derived from the special requirements, challenges, and scale of medical research. The ACME team explained, 'The gap arises in part from the specialised requirements of medical research …; in part from the special impetus and financial support enjoyed by medical research, which after all conveys the purpose of a distinctive NIH program; in part from the greater flexibility with which a service to a more coherent group can be administered compared to one announced campus-wide'.[489]

Convinced of the ACME's need and backed by Stanford's upper-level administration, Lederberg submitted the NIH proposal in October 1965, requesting $522,334 to purchase an IBM 360/50 computer and hire personnel to support this nascent biomedical computing operation. Not long after submitting the application, Lederberg received word of its approval. Once created, the system's day-to-day administration fell to Gio Wiederhold, who helped foster the ACME's growth, in the words of a 1968 annual report, 'from a primitive calculator system to one of the most powerful timesharing systems operating today'.[490] By 1968, the ACME was supporting a sundry of biomedical computing projects across the campus that ranged from the basic to the applied sciences—from 'the analysis and interpretation of data on human white cell antigens' to the processing of data on the 'effects of heart transplantation'.[491]

---

[489] Lederberg, 'Advanced Computer for Medical Research', pp. 3-4.
[490] 'Advanced Computer for MEdical Research (ACME)', Grant No. FR 00311-02, Special Research Resource Annual Report, 1 Aug. 1967-31 Jul. 1968, U.S. Department of Health, Education, and Welfare, Public Health Service, p. 3, JLP/60/29.
[491] 'Advanced Computer for Medical Research (ACME)', Grant No. FR 00311-03, Special Research Resource Annual Report, 1 Aug. 1968-31 Jul. 1969, U.S. Department of Health, Education, and Welfare, Public Health Service, pp. 91, 105, JLP/60/30-31.

Also among the computing projects supported by the ACME was one of the world's first expert systems, DENDRAL (DENDRitic ALgorithm). Expert systems are computer systems designed to emulate the decision making processes of human experts. They comprise two main subsystems: a knowledge base—a store of facts and rules about a given problem domain—and an inference engine that performs the reasoning process with respect to these rules and facts. DENDRAL's problem domain was organic chemistry; the system aimed to assist organic chemists in identifying the structures of amino acids and new organic molecules from mass spectrometry data. Although DENDRAL came to be supported by the ACME, the two enterprises were in fact homologous—they shared a common evolutionary lineage. Lederberg initiated the DENDRAL project (or at least its early antecedents) in 1961 as part of his on-going research into extra-terrestrial life ('exobiology'). At the IRL, Lederberg was developing an automated miniature laboratory for the Viking I Mars mission that would test for signs of life on Mars. DENDRAL, he hoped, would provide an important feature to this automated laboratory: using mass spectrometry data derived from a sample of Martian soil, DENDRAL could help identify amino acids and organic molecules that may be suggestive of extra-terrestrial life. Lederberg elaborated upon this vision in a draft proposal co-written with John McCarthy. They wrote, trying to convince NASA of its interest in supporting 'basic research' in artificial intelligence research, '[W]e believe that an automated biological laboratory in which a computer controls experimental apparatus according to programs that can be revised on earth is the key to the biological study of Mars'.[492]

[492] John McCarthy and Joshua Lederberg, 'A Proposal for the Study of Computer Control of External Devices and an Automated Biological Laboartory', draft proposal, 15 Oct. 1964, p. 2, JLP/67/69.

Though in many ways a product of Lederberg's vision, DENDRAL also benefited more circumstantially from its association with Stanford's intellectually rich environment. As Stanford grew through the 1960s into a massive research institution, countless prominent researchers across the life, physical, and medical sciences yielded to its gravitational pull. Among them was the eminent organic chemist Carl Djerassi, known for his role in inventing the oral contraceptive pill. In May of 1959, Djerassi wrote to Lederberg 'as a potential colleague', announcing his scheduled 'move to Stanford in September 1959'.[493] Before long, the two had developed 'a close personal and professional association', an association that quickly proved mutually beneficial.[494] As Lederberg would later explain, 'He was an accomplished mass spectrometrist; and of course I leaned very heavily on him for the elaboration of this technology for space applications. Conversely, he knew nothing about computers, and I was eager to find helpful applications in the zone of our common interest'.[495] Djerassi thus helped with the IRL's work using computers to address the 'data management problems of real-time gas-chromatography mass-spectrometry'—work that eventually evolved into DENDRAL.

At the same time that Lederberg was drawing upon Stanford's expertise in chemistry, he was also tapping into the surge of intellectual activity based at the university's young Department of Computer Science. On 6 April 1963, Lederberg attended a meeting organised by the Center for Advanced Studies in Behaviour Sciences on computer models of human thought. There, he met the young Edward Feigenbaum. A scholar who approached the world with great intellectual intensity, Feigenbaum had trained in computer science and artificial intelligence under the

[493] Carl Djerassi to Lederberg, 5 May 1959, JLP/9/202.
[494] Lederberg, 'How DENDRAL Was Conceived and Born', p. 6.
[495] Ibid., p. 6.

supervision of Herbert Simon, receiving his Ph.D. in 1960. Over the following five years, he grew interested in constructing models of empirical induction in science, and he was on the lookout for a concrete task domain within which to develop and test such models. Having been introduced to Lederberg in 1963, Feigenbaum felt comfortable, upon his move to Stanford in early 1965, to reach out. Feigenbaum asked the Nobel laureate if he knew of a concrete task domain 'that contained the essence of the empirical induction problem'.[496] In Lederberg's eyes, the problem of mass spectral analysis fit the bill perfectly: organic chemistry was a field that was not only rich in experimental data but also understood with respect to an accepted 'conceptual framework of mechanism'.[497] Such features made the field conducive to formalisation and computer modelling. Indeed, by 1965, Lederberg had already developed an algorithm, suitable for computer-oriented analysis, for calculating functional groups of organic molecules using high-resolution mass spectrometry data.[498]

Soon after linking up with Lederberg, Feigenbaum recruited the philosopher of science Bruce Buchanan, with whom he had worked for a summer at the RAND Corporation, to join the effort. Though trained as a humanist, Buchanan's background suited the DENDRAL project perfectly: in line with DENDRAL's early formulation as chiefly a project of scientific theory formation, Buchanan's academic expertise concerned the philosophy of scientific discovery. His doctoral dissertation advanced the proposition that discovery and theory formation in the sciences involved more of a systematic, formalizable process than had previously been recognised. 'I had some ideas', Buchanan explained, 'that it shouldn't be as mysterious a process as the

[496] Ibid., p. 9.
[497] Ibid., p. 14.
[498] Lederberg, 'Tables and an Algorithm for Calculating Functional Groups of Organic Molecules in High Resolution Mass Spectrometry', STAR (Scientific and Technical Aerospace Reports), No. N64-21426, 1 Apr. 1964, JLP/82/10.

psychologists were making it out to be with a creative, "Ah, ha, Eureka!" experience. It was totally mysterious. So I was trying to be somewhat systematic in looking at the process of discovery, but not from a psychological point of view, more from a logical point of view'.[499] Many of the academic publications concerning DENDRAL described the system precisely in this vein: 'More important', one such publication noted, speaking of the system's value to science, 'is the contribution such computer assistance can make to scientific creativity…. The automated approaches discussed in this paper provide a systematic procedure for verifying hypotheses about chemical structure and ensuring that no plausible alternatives have been overlooked'.[500]

A computer project with a narrow purpose—to assist organic chemists in identifying chemical structures—thus drew upon a wide spread of disciplines. It brought together an unlikely, and highly interdisciplinary, team comprising a geneticist, a computer scientist, an organic chemist, and a philosopher, among others. The interdisciplinarity and teamwork behind DENDRAL engendered an atmosphere of great excitement—albeit one that was male-dominated. Lederberg later singled out 'the fraternity that came out of the DENDRAL effort' as 'a high in [his] life experience, matching the gratifications of scientific excitement and … recognition'.[501]

The excitement translated into a working system. DENDRAL functioned in two steps. First, the system used an algorithm in combination with mass spectrometry data to enumerate all topologically unique configurations of a given set of atoms that also followed the requirements of chemical valence. Second, to narrow this list of all *possible* configurations down to a list of all *plausible* configurations, the programme

---

[499] Arthur L. Norberg, 'An Interview with Bruce G. Buchanan', Oral History 230, 11-12 Jun. 1991, Charles Babbage Institute, University of Minnesota, p. 4.
[500] Buchanan and Feigenbaum, 'Dendral and Meta-Dendral: Their Applications Dimension', p. 7.
[501] Lederberg, 'How DENDRAL Was Conceived and Born', p. 13.

then applied rules of thumb—a 'knowledge base' of heuristics—to assess the chemical plausibility of each candidate structure. Creating these rules was, in Lederberg's view, akin to 'cramm[ing]' the programme with chemical information. It was in the identification of such chemical information that Djerassi and his associates made distinctive contributions to the project, serving as 'founts of chemical expertise'.[502]

DENDRAL's reliance on domain-specific rules—that is, knowledge specific to the field of organic chemistry—heralded a new approach to artificial intelligence. Up to this point, most work in AI involved attempts to computerise and formalise some form of *general* human intelligence. Implicit in DENDRAL, however, was the view that human intelligence was little more than the domain-specific knowledge on which it rested. Lederberg recalled a revealing conversation with Feigenbaum on this very point—'a conversation when I was pressing Ed about the limitations of DENDRAL as general intelligence: he responded with the illumination that I may paraphrase: "That's exactly the point! Knowledge, not tricks or metaphysical insight, is what makes the program effective—and that itself is an insight of general import"'.[503] If there was any slogan under which DENDRAL travelled, it was this: 'Knowledge is power'.

MAKING MYCIN

In the fall of 1966, as Lederberg and his colleagues were working away on DENDRAL, a young student from Connecticut, Edward Shortliffe, moved to Cambridge, Massachusetts to begin his freshman year at Harvard College. Shortliffe gravitated toward Harvard's applied mathematics concentration (major), which was the main concentration at the time that contained a strong computer science bent. Owing to

[502] Ibid., p. 11.
[503] Ibid., p. 12.

the guidance of a mentor, Shortliffe discovered the recently-founded Laboratory of Computer Science (LCS) at the Harvard-affiliated Massachusetts General Hospital (MGH), where he first encountered the fields of medical informatics and biomedical computing. Directed by Octo Barnett, the LCS supported many lines of research at intersections of computing and medicine: the lab did work on electronic health records, decision support, programming languages, medical education, and medical information systems (Figure 23). At the LCS, Shortliffe became involved with research that was headed by one of Barnett's doctoral students, Robert Greenes. Greenes was developing, as part of his dissertation research, a computer-based system for medical record keeping, and Shortliffe brought, as Greenes remembered it, 'a high degree of technical skill, and … boundless energy and enthusiasm' to one area of the larger research project.[504] This area, which would become the topic of Shortliffe's own honours senior thesis, concerned the implementation of 'a structured input interface for composing research retrieval queries that could be run against the database of patients collected in the clinical system'.[505]

As his Harvard career came to a close, Shortliffe began looking to the future. Interested in combining a clinical career in medicine with a research career in biomedical informatics, he decided to apply for MD-PhD programmes, funded by the NIH through its Medical Science Training Program (MSTP), and commence training for a career as a physician-scientist. On learning of Shortliffe's wish to pursue a joint

[504] Robert A. Greenes, 'A Computer-Based System for Medical Record-Keeping by Physicians', Ph.D. Dissertation (Cambridge, MA: Harvard University, 1969). Robert A. Greenes, Bruce G. Buchanan, and Donald Ellison, 'Presentation of the 2006 Morris F. Collen Award to Edward H. (Ted) Shortliffe', *Journal of the American Medical Informatics Association* 14, no. 3 (2007): 376-385, p. 376.
[505] Greenes, Buchanan, and Ellison, 'Presentation of the 2006 Morris F. Collen Award', p. 376. Edward H. Shortliffe, 'A Computer-Based Medical Record System: Approaches to the Interface Problem', A.B. Thesis (Cambridge, MA: Harvard University, 1970).

career at the nexus of medical informatics and clinical medicine, Barnett steered him to Stanford. There, Shortliffe could take advantage of not just its strong programmes in computer science and medicine but also its more laissez faire research culture. 'He decided he wanted to go to med school', Barnett later recalled, 'and I told him to go to Stanford. At Stanford, he could get much more freedom to do what he wanted to do than [he could at] Harvard at the time'.[506]

Shortliffe arrived in Palo Alto in 1970. Having dipped his toe into a career combining computers and medicine while at Harvard, once at Stanford he jumped in with both feet and completed his joint degree in only six years. He spent the first and last two years working towards the MD portion, and he allotted himself an astonishingly brisk two years in the middle to complete his PhD. His actual habits of work, however, did not perfectly obey such demarcations; even fresh into medical school, he was thinking about computer science, artificial intelligence, informatics, and the applications of these fields to medicine. Not long after arriving at Stanford, he set out to master contemporary research in artificial intelligence, enrolling in a course on the Introduction to Artificial Intelligence and joining the Heuristic Programming Project (HPP) laboratory, which had grown directly out of the DENDRAL project and supported research into expert systems (of which DENDRAL was the oldest and most significant). In these environments, Shortliffe learned much through osmosis—through his proximity to exciting research into expert systems and to great minds like Lederberg and Feigenbaum. But Shortliffe was also, more than most dawning young medical students, the active author of his own educational and research programme.

---

[506] Octo Barnett, 'From "Farm Boy" to Director of the Laboratory of Computer Science: 2004 Interview of Octo Barnett', interviewed by Joan S. Ash and Dean F. Sittig, in Rebecca M. Goodwin, Joan S. Ash, and Dean F. Sittig (eds.), *Conversations with Medical Informatics Pioneers: An Oral History Collection* (Bethesda: U.S. National Library of Medicine, 2015), p. 10.

Observing that there was no regular journal club at the HPP, Shortliffe established one himself. The journal club, initially called 'SIGDoc', successfully brought together researchers at Stanford who were interested in the latest applications of computers to medicine.

It was through SIGDoc that Shortliffe, in 1971, met Bruce Buchanan, who would become his closest collaborator and one of his primary dissertation advisors (Figure 24). 'I quickly identified Bruce as an ideal colleague and mentor. He was thoughtful, provocative, experienced, and highly interested in applying DENDRAL-like ideas to a problem with greater clinical emphasis', Shortliffe recollected. 'Bruce was my closest advisor throughout this work and I still reflect fondly on the many hours that we spent in his office, working at the blackboard and devising new strategies for computational representation and control in that system'.[507] His other main advisor was the geneticist and clinical pharmacologist Stanley N. Cohen, whom Shortliffe had asked to work with upon his arrival at Stanford.

While Cohen would later become famous primarily for his research into bacterial plasmids and his role in developing recombinant DNA methodologies, at the time he was a natural pick to be Shortliffe's primary advisor. In the early 1970s, Cohen was effectively pursuing two careers, one as a basic scientist researching bacterial plasmids and another as a clinical pharmacologist exploring the applications of computers to clinical pharmacology. As Chief of the Department of Medicine's Division of Clinical Pharmacology, Cohen had been using the ACME to develop a computer-based reporting system that could monitor hospital drug use and alert doctors

[507] Shortliffe, to Shi-Kuo Chang, 5 Dec. 1990, Edward H. Shortliffe Papers, National Library of Medicine, National Institutes of Health (NIH), Bethesda, Maryland, box 2, folder 'Buchanan, Bruce, 1975-94' (hereafter cited in the form ESP/2/'Buchanan, Bruce, 1975-94').

of possible drug interactions. The system utilised a database of toxic drug interactions that Stanford's Division of Clinical Pharmacology had culled, evaluated, and synthesised from the published clinical literature over the course of two years. The contents of this database included extensive information about countless drug interactions: their pharmacological effects, mechanisms, anticipated symptomological presentations, clinical significance (i.e., the dosages at which they would be present), possible remediation, and bibliographic sources. With this interaction data, together with patient and prescription data entered by the pharmacist using cathode ray tube terminals, the system then used 'a highly interactive computer program to provide rapid notification to pharmacists, nursing staff, and physicians when a potentially interacting drug combination is prescribed'.[508] Cohen's system conveyed its findings to clinical staff through several different classes of reports that varied depending on the severity of the identified interaction. In light of his longstanding interest in using computers to improve clinical care, Cohen was eager to take on a doctoral student like Shortliffe.

With his advising structure in place, Shortliffe began to explore and define the parameters of a possible dissertation project. Given his interest in clinical medicine and his association with DENDRAL and the HPP, the possibility of bringing a DENDRAL-like system to the realm of medicine presented itself as the most natural path forward. Buchanan expounded, 'We thought that clinical medicine was ripe for the technology, and Ted believed it and took the ball and ran with it'.[509] Although MYCIN drew on DENDRAL, the two teams approached their respective projects in divergent ways.

---

[508] Stanley N. Cohen, Marsha F. Armstrong, Linda Crouse, and Gilbert S. Hunn, 'A Computer-Based System for Prospective Detection and Prevention of Drug Interactions', *Drug Information Journal* 72 (1972): 81-86, on p. 81. During this period, Cohen also published an influential handbook on drug interactions. Stanely N. Cohen, *Drug Interactions: A Handbook for Clinical Use* (Baltimore: Williams and Wilkins, 1974).

[509] Norberg, 'An Interview with Bruce G. Buchanan', p. 27.

DENDRAL, conceived by an eminent and polymathic scientist, was a top-down project—always moving with respect to the fixed North Star of Lederberg's grand vision. MYCIN, by contrast, was more of a bottom-up enterprise. To be sure, Shortliffe was brilliant, driven, and resourceful, and he wanted to explore computer applications in medicine. But within this broad domain, he was less prescriptive about the particulars: he was happy to go wherever his research or circumstance took him. As a result, the aims of the project that came to be known as MYCIN shifted over the course of its development.

Shortliffe, Buchanan, and Cohen initially planned to create a programme that would evaluate decisions that had already been made at a hospital and flag the subset of those that were problematic and in need of peer review. The programme, in other words, was to help clinical practice review boards *retrospectively* determine when physicians had erred in their prescribing. As the team began to work on the programme, however, they quickly decided that a forward-looking *prospective* programme, one that offered advice about yet-unmade medical decisions, might be more productive, beneficial, and acceptable than the backwards-looking retrospective programme that they had originally envisioned. Why run around stomping out forest fires when there was an opportunity to prevent them from starting in the first place? Buchanan elaborated on this shift in an interview: '[W]e decided that critiquing was not the right psychological model for a computer program with physicians' actions. But prospective systems to keep people out of trouble might be more acceptable'.[510]

As an expert system that helped diagnose and treat microbial infections, the MYCIN project bridged the interests of Shortliffe's two main advisors. It combined Buchanan's intimate knowledge of artificial intelligence and expert systems with

[510] Ibid.

Cohen's expertise in clinical pharmacology and antimicrobial resistance. Cohen would later describe how his engagement with the MYCIN project 'meshed well with [his] interest in resistance to antimicrobial drugs'.[511] Having operationalised the project—having identified a problem domain and an approach to making decisions within that domain—now the team had to bring the programme to life. Among the most urgent orders of business was to start building MYCIN's knowledge base—its collection of rules about the domain of bacterial infections and antibiotic treatment. Such knowledge, however, fell outside the bounds of Shortliffe's, Buchanan's, or even Cohen's narrow area of expertise. 'An important aspect of this work', a 1973 MYCIN report pointed out, 'is its strong interdisciplinary nature—success in this project depends in part upon the integration of medical knowledge into a computation framework. This requires the participation of professionals with differing expertise, followed by the close coordination of the people from the various departments working with it…'.[512] If Shortliffe and his colleagues in the Department of Computer Science were going to get MYCIN to work, they were going to have to build connections with others at the School of Medicine, particularly those within the Division of Infectious Disease.

Luckily, Cohen already enjoyed personal and professional connections to Stanford's medical faculty. Exploiting this network, Cohen corralled Thomas Merigan, Chief of the Division of Infectious Disease, and Staton Axline, a physician in that division, into joining the project. (It was Axline who, in 1973, suggested that the system be named MYCIN.) By March 1972, a semi-formal, interdisciplinary group of

---

[511] Stanley N. Cohen, 'Science, Biotechnology, and Recombinant DNA: A Personal History', oral history conducted by Sally Smith Hughes in 1995, Regional Oral History Office, Bancroft Library, University of California at Berkeley, 2009, p. 23.
[512] 'Computer Acquisition of Medical Judgmental Knowledge', research proposal, Dec. 1973, ESP/10/'MYCIN History, 1972-78'.

researchers had coalesced around the MYCIN project. The group met on a weekly basis to discuss all aspects of the MYCIN system, from the nature of decision making in the context of antibiotic therapy, to more technical aspects of the system's programming and operation. Present at the initial meetings were Shortliffe, Buchanan, Cohen, Axline, Merigan, and the clinical pharmacologist G. Hunn. Cordell Green—a computer science professor who taught the course at Stanford from which Shortliffe had learned LISP programming (CS 206: Symbolic Expressions)—also joined the meetings at a later time.

An early plan for the antibiotic therapy programme was for it to build directly upon existing computer-based operations at the School of Medicine. Perhaps, the group speculated, the programme could simply extend the clinical pharmacology programme that had already been developed by Cohen. Using organism and resistance pattern information obtained from the clinical microbiology laboratory, the MYCIN group wondered if Cohen's pharmacology programme could be made to advise physicians on the appropriateness of antimicrobial therapy in addition to warning them about drug interactions. The team, however, soon came to appreciate that decision making with respect to antibiotic therapy required far more than data from the clinical microbiology laboratory alone. As Shortliffe related in a 1972 memo and progress report, 'It quickly became apparent … that the decision model used by an expert on antimicrobial therapy before he recommends treatment is highly complex and requires answers to questions only a few of which can be answered by the clinical microbiology lab culture and resistance pattern data'.[513]

[513] Edward Shortliffe to Stanley N. Cohen (S. Axline, B. Brown, B. Buchanan, C. Green, G. Hunn, T. Merigan, B. Raphael, J. Rulifson, R. Schimke, P. Suppes, H. Sussman, and R. Waldinger copied), 'Progress Report on the Antimicrobial Therapy Project', memorandum, 13 Oct. 1972, ESP/10/'MYCIN History, 1972-78'.

What, then, did antimicrobial therapy selection involve? This question fell to the group's infectious disease experts, chiefly Stanton Axline. In February 1972, Cohen sent a memo to Merigan and Axline on the subject of their '[p]roposed computer project involving antibiotics therapy'.[514] The memo summarised the early plans for MYCIN and outlined the role Merigan and Axline, as sources of human expertise, would play in the system's creation. The two infectious disease specialists would, the memo noted, play a central role in defining and identifying the kinds of knowledge that physicians use when identifying and treating infections: 'It was agreed that Tom [Merigan] and Stan (Axline) would compile a list of major clinical factors to be considered in preparing the initial program for acquisition of the data base (i.e. the major branch point in the decision making process re: prescribing a specific antibiotic for a given situation)'. Within a few weeks, Axline responded to this charge. In a letter to the members of the original MYCIN group, Axline enclosed 'a preliminary list of major clinical factors used in the decision-making process for antibiotic therapy of bacteremia'.[515] He enumerated the principal clinical factors that physicians draw upon in making their decision—factors that came not only from the bacteriology laboratory (e.g., Gram stain result) but also the pharmacy (e.g., drugs being taken at the time of culture collection), the patient's chart (e.g., the presence of catheters in place in any cavity), and other domains of medicine.

Through the spring of 1972, the MYCIN team explored other existing systems off of which they might model MYCIN. Shortliffe and Buchanan met once or twice

---

[514] Stanley N. Cohen to Thomas Merigan and Stanton Axline, 'Proposed Computer Project Involving Antibiotics Therapy', memorandum, 25 Feb. 1972, ESP/10/'MYCIN History, 1972-78'.
[515] Stan Axline to Cohen, Tom Merigan, Buchanan, Gil Hunn, Shortliffe, 'Proposed Computer Project Involving Antibiotics Therapy', memorandum, 13 Mar. 1972, ESP/10/'MYCIN History, 1972-78'.

each week to study one such programme—SCHOLAR, which had been developed by the electrical engineer Jamie Carbonell at MIT. A tutoring programme, SCHOLAR used semantic networks derived from AI research to help represent knowledge about South American geography in a way that was conducive to computer-aided instruction.[516] After working with SCHOLAR for a few months, however, Shortliffe and Buchanan threw in the towel: they concluded that 'the SCHOLAR approach was inappropriate for the kind of problem area with which [they] were concerned'.[517] The heart of the problem was the sharp difference between the world of medicine and the world of geography. They found the domain of infectious disease to be too ill-defined and ill-structured for representation by semantic network; infectious disease did not enjoy the same clear boundaries, borders, and demarcations that were characteristic of SCHOLAR's domain of South American geography. By June the team had arrived at a different strategy. They garnered 'an improved picture of the problems inherent in our task and decided that a modified theorem-proving approach based upon situation-action rules and multiple goals was the best way for us to proceed'.[518]

MYCIN thus emerged as a 'rule-based' expert system. This meant that MYCIN's knowledge base consisted of rules, hundreds of 'IF-THEN', 'premise-clause' statements about the diagnosis and treatment of bacterial infections. Consider rule 047:

IF: 1) The site of the culture is blood, and
2) The identity of the organism is not known with certainty, and
3) The stain of the organism is gramneg [gram negative], and
4) The morphology of the organism is rod, and
5) The patient has been seriously burned

[516] Jamie R. Carbonell, 'AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction', *IEEE Transactions on Man-Machine Systems* 11, no. 4 (1970): 190-202.
[517] Shortliffe to Cohen, 'Progress Report on the Antimicrobial Therapy Project', 13 Oct. 1972.
[518] Ibid.

THEN: There is weakly suggestive evidence (.4) that the identity of the organism is pseudomonas.[519]

MYCIN used its collection of rules like this to reason through medical cases using a process called 'backwards chaining'. In forward chaining, computerised systems start with the set of facts about a case, only later bringing the relevant rules to bear on these facts. Backwards chaining goes in reverse: instead of starting with the facts at hand, the system begins with a goal—a piece of information to find out, a diagnosis to make, or some other end point—and then invokes the rules that can help the computer reach that goal. For example, MYCIN may proceed with the hypothesis, or 'goal', that the 'identity of the organism is pseudomonas'. The system would identify the rules that assist in making this conclusion. Each of these rules, such as rule 047 above, would have a series of premises (e.g., 'the morphology of the organism is rod') that need confirmation and that other rules may address. The system thus proceeds backwards, reaching back to these other rules—or reaching out to the physician—in order to establish the premise clauses of this final rule. MYCIN essentially worked through various hypotheses or goals in this fashion until it reached its conclusions about the case.

The initial batch of rules incorporated into MYCIN came from the team's infectious disease specialists, predominantly Axline and Cohen. The team presented Axline and Cohen with actual patient infections as represented in patient charts. Before the team's computer scientists, Axline and Cohen performed their diagnostic and therapeutic reasoning with respect to these real cases, attempting to explicate their process of decision making and distil their knowledge into discrete rules. 'We reviewed several patient records', Shortliffe noted in 1972, 'and began to examine what kinds of

[519] Shortliffe, *Computer-Based Medical Consultations:* MYCIN, p. 52.

rules the experts in our group were using in an attempt to reach a therapeutic decision'.[520] He elaborated on this process of rule extraction and creation during a 1975 workshop:

> All the rules we have in our system have been acquired at weekly meeting in which Dr. Axline and Dr. Cohen, two clinicians most closely associated with our project, took patient charts and with the end of those charts still unknown to them, began to review them. Those of us unfamiliar with the clinical aspects of what was being discussed would listen and try to pick out the underlying threads of reasoning. We would then code these into rules and use them to run patients' charts.[521]

Once formulated, the rules were shopped back to the medical experts, who judged the extent to which these preliminary rules were understandable and persuasive: 'We would then bring back the results to show the expert how the system actually used the rules in order to come up with recommendations. So our concerns were whether or not the rule we used represented a fact that the expert could agree with'.[522] In other words, MYCIN's production rules, if not crafted with the direct intent of imitating human reasoning, were crafted with an interest in the discernibility of these rules to human reasoning.

The development of MYCIN therefore depended on the assumption that medical knowledge was *communicable*—that it could be extracted from medical experts, crystallised into discrete packets of information, and programmed into a computer system (Figure 25). Medical knowledge concerning antibiotic therapy was presumed to be, for the infectious disease specialists, explainable and, for the computer scientists, comprehensible. Certain *social* properties of the MYCIN project during these

---

[520] Shortliffe to Cohen, 'Progress Report on the Antimicrobial Therapy Project', 13 Oct. 1972.

[521] Amarel, Kulikowski, Sridharan, and Sridharan, *Proceedings of the First Annual AIM Workshop*, p. 23.

[522] Ibid., p. 24. Emphasis added.

early stages facilitated the program's development. The MYCIN team was fairly intimate and small: every member knew every other member, and they all met regularly to touch base, trade ideas, and update the system. Second, at the early stages of development, the MYCIN system itself, like the team that created it, was also pretty 'small'. By October 1972, MYCIN contained just 70 rules.[523] Thanks to the group's intimacy and the system's workability, a tight feedback loop formed between those who explicated the rules and those who encoded and tested them.

REASONING UNDER UNCERTAINTY

Medicine is an uncertain business. It is, William Osler famously mused, 'a science of uncertainty and an art of probability'.[524] Doctors must make medical decisions with limited knowledge and imperfect data, and they must make these decisions with respect to individual patients who sometimes respond to equivalent treatments in surprising, divergent, and unpredictable ways. That a given therapy works in 95 per cent of patients is no guarantee that it will work in *this* patient. As the sociologist Rene Fox has noted, a central part of becoming a doctor is learning to live and cope with the limits of medical knowledge, the vagaries of individual patients, and one's own personal experiences of puzzlement, doubt, and ambiguity; training to become a doctor is often a matter of 'training for uncertainty'.[525] As they did medical trainees, the problems of uncertainty in medicine also plagued the MYCIN team. No

---

[523] Shortliffe to Cohen, 'Progress Report on the Antimicrobial Therapy Project', 13 Oct. 1972.
[524] William Osler in W.B. Bean (ed.), *Aphorisms from His Bedside Teachings and Writings*, ed. 3 (Springfield, IL, 1961), p. 129.
[525] Rene Fox, 'Training for Uncertainty', in Robert K. Merton, G. Reader, and P. L. Kendall (eds.), *The Student Physician* (Cambridge, MA: Harvard University Press, 1957). See also Rene Fox, 'The Evolution of Medical Uncertainty', *The Milbank Memorial Fund Quarterly* 58, no. 1 (1980): 1-49. For a historical perspective on uncertainty in medical decision making, see Jones, *Broken Hearts.*

aspect of developing MYCIN consumed more thought and energy than the question of how to formalise such uncertainty—how to convey, represent, and determine the strength of belief about a medical rule or proposition that is not known with complete certainty.

Bayes' theorem offered one way forward. By the 1960s, many research teams had embraced Bayesian approaches to the task of computerising diagnosis. The HEME project adopted such an approach, as did Shortliffe's own undergraduate mentor. In 1967, Octo Barnett, together with Anthony Gorry, developed an adapted Bayesian model of diagnosis for computer implementation.[526] Barnett and Gorry's diagnostic model developed upon a more prominent Bayesian system that had come out of the University of Utah a few years earlier (see Part II). There, Homer Warner and colleagues created a programme for the diagnosis of congenital heart disease.[527] For all their promise, Bayesian approaches to computerising diagnosis employed certain assumptions that rarely held in the context of clinical medicine. For example, Bayes' theorem assumes that probability inputs are conditionally independent, but in medicine this is rarely the case: diagnostic variables tend to change together, with a given disease state affecting multiple variables at once.

---

[526] G. Anthony Gorry and G. Octo Barnett, 'Experience with a Model of Sequential Diagnosis', *Computers and Biomedical Research* 1 (1968): 490-507.
[527] Homer R. Warner, A.F. Toronto, and L.G. Veasy, 'Experience with Bayes' Theorem for Computer Diagnosis of Congenital Heart Disease', *Annals of the New York Academy of Science* 115 (1964): 558-567. For other examples of Bayesian programmes during this period, see G.S. Lodwick, C.L. Haun, W.E. Smith, R.F. Keller, and E.D. Robertson, 'Computer Diagnosis of Primary Bone Tumors: A Preliminary Report', *Radiology* 80 (1963); V.X. Gledhill, J.D. Mathews, and I.R. Mackay, 'Computer-Aided Diagnosis: A Study of Bronchitis', *Methods of Information in Medicine* 11 (1972): 228-233; R.P. Knill-Jones, R.B. Stern, D.H. Grimes, J.D. Maxwell, R.P.H. Thompson, and R. Williams, 'Use of a Sequential Bayesian Model in Diagnosis of Jaundice by Computer', *British Medical Journal* 1 (1973): 530-533.

More problematic than the assumptions required for a functional Bayesian diagnostic system were its data requirements. For a Bayesian system to produce accurate results, it requires large amounts of data—data that allows for the computation of the conditional probabilities on which the theorem rests. This data didn't just have to be big; it also had to be good: it had to reflect the true statistical properties of the disease in question. Warner and his colleagues in Utah possessed such a data set, with high-quality information on hundreds of patients with congenital heart defects. But Shortliffe—and countless other physicians—were not so fortunate: these clinicians were working in areas of medicine where clinical judgement, not hard data, was king. Shortliffe and his colleagues were thus presented with a conundrum: how, exactly, to formalise medical reasoning when that reasoning is itself *inexact*? As the problem of poor-quality data was ubiquitous in medicine, Shortliffe hoped that an alternative to the Bayesian approach might have wide appeal. 'Although conditional probability provides useful results in areas of medical decision making such as those I have mentioned, vast portions of medical knowledge suffer from so little data and so much imperfect knowledge that a rigorous probabilistic analysis, the ideal standard by which to judge the rationality of a physician's decisions, is not possible'.[528]

These weren't just different approaches to the computerisation of diagnosis; they were also scaffolds around which the groups constructed different professional identities. Members of Warner's group of the University of Utah viewed the efforts at Stanford as overly academic, theoretical, detached.[529] The Utah team thus projected themselves as Stanford's foil: pragmatic, data-driven, applied. Warner later reflected on

---

[528] Edward H. Shortliffe, *Computer-Based Medical Consultations: MYCIN* (New York: Elsevier, 1976), p. 160.

[529] This dynamic is one that played out more widely across the fields of mathematics, statistics, and computer science, particular in the post-war period. See, e.g., Jones, 'How We Became Instrumentalists (Again)'.

this dynamic in an interview: 'I'd always been sort of a critic of artificial intelligence, and the whole notion of it. I think it's far-fetched….I think Ted's [Shortliffe] done some good things, even though they're more on the theoretical level, and we've been over on the practical end'.[530]

In 1975, Shortliffe and Buchanan first described their solution to the problem of uncertainty in the landmark paper, 'A Model of Inexact Reasoning in Medicine'.[531] The uncertainty framework that they proposed and implemented rested on the concept of the certainty factor (CF). The certainty factor—a value that ranged from -1 to +1—reflected the change in belief in a hypothesis given certain evidence. Compared to Bayesian analysis and conditional probabilities, the CF concept, Shortliffe and Buchanan believed, was more intuitively accessible to physicians, who often had difficulties accurately estimating probabilities in the way a Bayesian system would require but who were quite comfortable expressing general 'degrees of belief' about their medical conclusion. CFs, the team argued, captured and operationalised these degrees of belief.

Shortliffe and Buchanan were—despite how Warner may have framed it—pragmatic: they would not let the perfect be the enemy of the good. Their primary aim was producing a model that would work. When, in 1975, an MIT graduate student wrote to Shortliffe criticising aspects of the CF model, Shortliffe defended his model not as an unassailable finished product but as a practical step in the right direction. 'I guess you know that our own reaction to the CF model is based upon pragmatics and

---

[530] Homer R. Warner, 'Founder of the HELP System and the Utah Medical Informatics Program', interviewed by Dean F. Sittig, in Rebecca M. Goodwin, Joan S. Ash, and Dean F. Sittig (eds.), *Conversations with Medical Informatics Pioneers: An Oral History Collection* (Bethesda: U.S. National Library of Medicine, 2015), p. 20.
[531] Edward H. Shortliffe and Bruce G. Buchanan, 'A Model of Inexact Reasoning in Medicine', *Mathematical Biosciences* 23 (1975): 351-379.

the need to have something that seems to work and to be at least mildly defensible within the domain for which it is designed'.[532] Shortliffe and his colleagues, having found themselves navigating a *terra incognita*, improvised their way forward. Shortliffe continued, '[T]here was really very little attempt in the literature to come up with an actual quantitative model. We were left to follow our instincts with the goal of getting together a version of MYCIN that would work acceptably well. If another model is proposed that improves upon problems inherent in the CF approach, we would switch over in a minute. We have no vested interest in the CF's per se; we just don't see any real alternatives around'.[533]

Their focus on the system's function did not mean the team was uninterested in its form. Even as the MYCIN team strove to create a system that worked, they also emphasised certain qualities of *how* MYCIN worked. MYCIN, they suggested, approximated—and gave a quantitative shape to—the 'artistic' and 'inexact' forms of reasoning used in medicine. A key contribution of MYCIN, the group argued, was to capture, formalise, and computerise this mode of reasoning that heretofore had seemed to defy quantitative analysis. They explained, 'The purpose of this paper is to examine the nature of such nonprobabilistic and unformalized reasoning processes and to propose a model by means of which such incomplete "artistic" knowledge might be quantified'.[534] MYCIN therefore was developed to 'reflect' (at least roughly) how experts reason with medical knowledge: '[W]e look for ways to handle decision rules as discrete packets of knowledge and for a quantification scheme that permits

[532] Shortliffe to William Mark, 21 Jul. 1975, ESP/10/'MYCIN, 1973-77'.
[533] Ibid.
[534] Shortliffe and Buchanan, 'A Model of Inexact Reasoning', p. 353.

accumulation of evidence in a manner that adequately reflects the reasoning process of an expert using the same or similar rules'.[535]

Claims about the proximity of computerised reasoning to human reasoning—that is, claims about man-machine fidelity—often shaded into claims about trust, transparency, and communicability. Models of reasoning that approximated modes of human reasoning, many argued, could be better understood by, and communicated to, human users.

OF MAINTENANCE AND MODULARITY

In Lewis Carroll's *Through the Looking-Glass*, the young Alice, a visitor to Looking-Glass Land, notices something odd while running. No matter how quickly she moved her feet, she realised that she wasn't moving with respect to the things around her. Perplexed, she bemoans to the Red Queen, 'Well, in our country…you'd generally get to somewhere else—if you run very fast for a long time, as we've been doing'.[536] The Red Queen interjects, explaining that the laws governing Looking-Glass Land were different from those governing Alice's country: 'Now, *here*, you see, it takes all the running you can do, to keep in the same place'. Because everything in Looking-Glass Land was in a perpetual state of motion, one needed to continually move just to stay put.

Developers of computerised diagnostic systems like Shortliffe and his colleagues discovered that the world of medicine shared much in common with Looking-Glass Land. The MYCIN team was not working in a domain where knowledge, information, and data were frozen or fixed in time; knowledge about

[535] Ibid., p. 356.
[536] Lewis Carroll, *Through the Looking-Glass* (Mineola, NY: Drover Publications, Inc., 1990), p. 20.

antibiotic prescribing, and about medicine in general, was growing and shifting in real time. Thus, even after versions of MYCIN had been put forward, the system had to account for the fact that its knowledge base—that very thing that gave the system its power—was in a state of on-going evolution and revision. The team was perpetually looking for ways to add new rules and tweak or improve the formulation of existing rules. To create the system and then step away would be to let the system fall into obsolescence—to let it fail.

Development of MYCIN was therefore less a singular act of invention than a continual process of maintenance. Historians of technology have long focused on innovation—on great individuals and their fleeting acts of creation and genius. But a recent turn within the historiography has shifted attention to the centrality of maintenance.[537] Andrew Russell and Lee Vinsel in particular have described and organised this recent push to 'study momentum and path dependence, use and users, repair and breakdown—those things that may be swept away in the Schumpeterian gale of creative destruction'.[538] New technologies do not come into being through one-off strokes of genius but through the unglamorous work of tweaking, updating, cleaning, adapting, improving, editing, debugging—over and over again. Getting technologies built—and subsequently getting them to work—often falls at the feet of 'maintainers'.

MYCIN bears these conclusions out. The system's knowledge base was under a constant state of construction. In 1973, the MYCIN team compared the growth and

---

[537] See, e.g., David Edgerton, *The Shock of the Old: Technology and Global History Since 1900* (Oxford and New York: Oxford University Press, 2007); Andrew L Russell and Lee Vinsel, 'Making Maintainers: Engineering Education and the Ethics of Care', in Matthew Wisnioski, Eric Hintz, and Marie Stettler Kleine (eds.), *Does America Need More Innovators?* (Cambridge, MA: MIT Press, 2019). Maintenance has also long been an enduring and vibrant theme of certain feminist histories. See, e.g., Susan Strasser, *Never Done: A History of American Housework* (New York: Pantheon, 1982).
[538] Andrew Russell and Lee Vinsel, 'After Innovation, Turn to Maintenance', *Technology and Culture* 59, no. 1 (2018): 1-25, p. 3.

education of MYCIN to that of a medical trainee. Far from 'complete', MYCIN was still in the process of acquiring medical knowledge and learning from its mistakes: 'The program is currently at a stage roughly analogous to that of a medical student who, having completed his coursework, enters the clinic and encounters real and challenging cases. Like the student, the program will only gain a high level of broad competence by an extended period of being advised of its shortcomings and corrected in its errors'.[539] A relatively young system, MYCIN still had much to learn. By encountering patient cases, real or contrived, in controlled environments, the gaps in MYCIN's knowledge could be identified and addressed.

The need to maintain and update the system was not only a consequence of MYCIN's relative youth and inexperience. It was also a side effect of the nature of the knowledge domain itself. To carry forward the team's analogy, even if MYCIN reached the level of an attending, the work of maintenance would not be done. Just as attendings must remain 'up to date', incorporating new knowledge and evidence into their patterns of care, MYCIN too would have to incorporate new information about antibiotic therapy—new bugs, new drugs, new toxicities, new patterns of susceptibility. '[N]either the CF model nor MYCIN', Shortliffe wrote in 1975, 'is static'.[540]

Beyond having to attend to MYCIN's knowledge base, the team also had to update and revise the system's code. The work of 'debugging' accounted for a significant chunk of the time expended towards MYCIN's development. Members of the team recalled 'very bad debugging experiences'—experiences of 'wading through code' in search of some underlying cause of the system's problems.[541] The team

---

[539] 'Computer Acquisition of Medical Judgmental Knowledge', research proposal, Dec. 1973.
[540] Shortliffe to Mark, 21 Jul. 1975.
[541] Clancey to Greg Cooper (Shortliffe and Buchanan copied), email printout 6 Oct. 1980, ESP/2/'Clancey, William (1), 1978, 81'.

reiterated these sentiments in a 1973 research proposal on 'Computer Acquisition of Medical Judgmental Knowledge': 'Current attempts to create knowledge-based programs have underscored the difficulties of such understandings—vast amounts of work have been required for the derivation, encoding, and debugging of the knowledge bases'.[542] Because code maintenance could easily burn through large amounts of time, the team created features that aimed to speed up the debugging process. MYCIN member Carlisle Scott described one of these features in a 1978 email: 'The main function of BATCH is for debugging the knowledge base. System builders can create a library of stored cases for use in testing the system'.[543]

Luckily, MYCIN was premised on an understanding of medical knowledge that seemed conducive to revision and change. A collection of discrete IF-THEN rules, new rules could be easily added to the soup, and old, bad rules easily removed. Medical knowledge, in other words, was seen to be modular, consisting of individual building blocks that together amounted to the vast edifice of medical knowledge. The modularity of rules, some reasoned, meant that MYCIN's knowledge base could be easily understood, explained, and vetted—its errors and biases easily rooted out. The proposed introduction to a 1988 conference on the topic of 'expert system knowledge transfer' adopted this view of rule-based expert systems like MYCIN:

> The clinical expertise and interpretive rules in the knowledge base can be reviewed, modified and expanded independently by clinicians who are not required to have programming capabilities. Any biases, judgements, or logical errors remain exposed for scrutiny and discussion by the clinical, and not only by the programming community. This approach also allows the breadth and power of the system to be improved by adding to and modifying the knowledge base in a modular fashion.[544]

---

[542] 'Computer Acquisition of Medical Judgmental Knowledge', research proposal, Dec. 1973.
[543] Carli Scott, 'BATCH', email printout, 13 Nov. 1978, ESP/11/'MYCIN – General, 1973-82'.
[544] Ove Wigertz to Shortliffe, email printout, 3 Oct. 1988, ESP/3/'Columbia Presbyterian Med Ctr Retreat, (Arden Homestead, 6/89), 1989'.

Early papers from the MYCIN group similarly listed 'ease of modification' as a key advantage of the rule-based approach over Bayesian statistical approaches: 'since the rules are not explicitly related to one another and there need be no pre-structured decision tree for such a system, rule modifications and the addition of new rules need not require complex considerations regarding interactions with the remainder of the system's knowledge'.[545]

This was the hope, at least. In reality things were more complex. Adding new rules to MYCIN's knowledge base sometimes produced unpredictable changes to the system's operation. This point was raised by the infectious disease fellow Robert Blum during a presentation at a MYCIN team meeting in September 1977. Blum's stance was summarised in the meeting minutes: 'Criticized rule-based systems on the basis of lack of modularity. The chunks are no[t] independent, statistically or otherwise. Pointed to problem associated with entry of new rules, i.e., that experts must be familiar with the entire corpus of rule[s] to be able to enter "independent" new ones'.[546]

Challenges that the team encountered in building and refining MYCIN supported Blum's critique of modularity. William Clancey, for example, documented what he called the 'rippling of CF's within the body of rules'.[547] Specifically, Clancey and others realised that, because MYCIN employed methods of combining certainty factors—of bringing several pieces of evidence to bear on a given hypothesis—unexpected interactions sometimes arose between rules and their associated certainty factors. 'As a result of this study', Clancey explained, 'rules 4 and 27 have had their

[545] Shortliffe and Buchanan, 'A Model of Inexact Reasoning in Medicine', p. 357.
[546] Robert Blum, 'Blum's presentation' email printout, 6 Oct. 1977, ESP/10/'MYCIN-Current, 1975-82'.
[547] Clancey, 'MYCIN's Sensitivity to CF changes', unpublished report, Sep. 1975, ESP/10/'Certainty Factor, 1976-84'.

CF's changed to .3 (instead of. 2). The moral here is that, while rules are supposed to be modular and rule CF's are to be chosen independently, the system will not function properly if the foundation-confirmation scheme is used and CF interactions aren't thought out carefully. It pays to keep in mind the entire body of rules which conclude a given parameter's value'.[548]

Assumptions about the modularity of medical knowledge also came under strain as the MYCIN team worked to combine different knowledge bases and rule sets. The original MYCIN knowledge base concerned the infectious disease domain of bacteraemia—bacterial infections of the blood. Soon, other members of the team—principally the infectious disease fellow Victor Yu—began developing a knowledge base that would allow MYCIN to operate in the area of acute meningitis as well. By the mid-late 1970s, the team was working to combine the two rules sets. Due to rule-rule interactions, however, combination of the bacteraemia and meningitis knowledge bases was a process more complex than one of mere addition. The team reported on these challenges in the minutes of a 15 June 1976 meeting: 'Larry [Fagan] and Carli [Scott] are starting on integration of both sets of rules into the new control structure. There will be some interactions of rules that will cause trouble'.[549]

Experiences like these led Clancey to declare in a March 1976 memo: 'For too long I have assumed that because a rule looks like a discrete object it is necessarily modular'.[550] In MYCIN, then, one sees the co-production of this rule-based expert system and a modular understanding of medical knowledge: the system arose out of

[548] Ibid.
[549] Carli Scott, 'MYCIN meeting June 7, 1978', email printout, ESP/10/'MYCIN-Current, 1975-82'.
[550] Clancey, unpublished memo, 1 Mar. 1976, ESP/10/'Certainty Factor, 1976-84'.

ideas about modularity—ideas that the system in turn reinforced through its development and operation.

Shortliffe was sympathetic to these sentiments. On reading the draft introduction of the proposed 1988 conference on 'expert system knowledge transfer', Shortliffe expressed surprise over the lack of critical discussion around modularity. He wrote to the conference organisers, 'The nitty gritty syntactic details will only be worth pursuing if we have a common view of the big picture. I might mention, for example, that some of my students recently wrote a paper entitled "the myth of modularity" which strikes at the very heart of the presumptions that knowledge can be created and stored in separable packets'.[551] While Shortliffe did not disregard claims of non-modularity, the view did strike him as 'too extreme'. Even if rules might occasionally interact in surprising ways, he felt that the modularity of rules could be defended—if only on a pragmatic basis. As Shortliffe wrote to his colleagues in 1976, 'I still tend to believe that the claims of rule modularity can be defended although there are clearly examples of rules that are poor precisely because they do not conform to this requirement'.[552]

The challenges around revising MYCIN's knowledge base and combining its bacteraemia and meningitis rules highlights how the social dynamics of the MYCIN team play into the system's development. By the late 1970s, countless individuals had

---

[551] See David Heckerman and Eric Horvitz, 'The Myth of Modularity in Rule-Based Systems for Reasoning with Uncertainty', *Machine Intelligence and Pattern Recognition* 2 (1986): 23-34. See also Eric Horvitz and David Heckerman, 'Modular Belief Updates and Confusion about Measures of Certainty in Artificial Intelligence Research', in Laveen Kanal and John Lemmer (eds.), *Proceedings of the First Conference on Uncertainty in Artificial Intelligence* (Arlington, VA: AUAI Press, 1985), 283-286.

[552] Shortliffe to Scott (Yu, Buchanan, David, van Melle, Clancey, Fagan, Axline, Wraith, and Aikins copied), email printout, 8 Jul. 1976, ESP/10/'MYCIN – Current, 1975-82'.

taken part in MYCIN's development, writing new rules, tweaking old rules, and debugging the code underlying all rules, new and old. Whereas one group of people was involved in developing MYCIN's bacteraemia knowledge base, another group—headed by Victor Yu—was involved in creating the meningitis knowledge base.

The methods and assumptions behind the bacteraemia and meningitis rules were not always aligned. Yu, for example, used smaller gradations between different certainty factor values for the meningitis rule set. Specifically, Yu selected certainty factors to the nearest 0.05. The decision sparked controversy among the MYCIN gang. Some questioned whether such precise gradations carried any real clinical meaning. Could a physician, or the state of medical knowledge more generally, truly defend that a given statement was known with a 0.45, as opposed to a 0.40, level of certainty? And if such a degree of precision could be defended, would these fine distinctions even affect clinical decision making in the real world? 'What remains disturbing is the certainty factor model itself', explained Clancey. 'Here we have no sure intuition about the performance meaning of .05 as opposed to .1, yet we are assigning them as if they were significantly different from one another'.[553]

Particularly contentious was the way in which Yu used certainty factor values below the 0.2 cutoff. Below the 0.2 value, evidence was deemed to be 'not known'. To Clancey, it didn't make sense to assign CF gradations below 0.2. 'How is it that "evidence" which is too weak to yield an acceptable conclusion nevertheless is definite enough to be put in one of the three CF categories: .05, .10, and .15?'[554] Moreover, Clancey worried that making CF assignments at these lower ranges might result in peculiar and nonsensical outcomes, with multiple CF values in the 'not known' range

[553] Clancey, '<CLANCEY>CF.TWOCUTOFF;1', email printout, ESP/10/'MYCIN History, 1972-78'.
[554] Clancey, unpublished memo, 1 Mar. 1976.

combining to produce an actionable finding or observation. As Clancey argued in a message to the group,

> I am saying here that rules which break observations into many small parts, resulting in CF's <.2 intended to combine to form an accumulated observation, ignore the total perspective which says, "Hey, wait a minute, these 6 clauses can't appear together: why was she given corticosteroids if she has XX? This doesn't mean FUNGUS to me; no, I want to know why that prescription was made." This same criticism does not apply with the same force to many rules with CF >.2 because they bring together a "more significant set of facts." They do this by capturing (often disjoint) pictures of the world which in themselves MAKE SENSE. I do not at all understand how a rule can be written which can at once stand on its own, and yet NOT be significant truth (i.e., believable observation, tangible conclusion).[555]

For all of the controversy around Yu's approach, the meningitis rules seemed to outperform the original bacteraemia knowledge base in at least one important respect: the meningitis rule set seemed to respond more favourably and predictably to subsequent additions and edits. The team summarised this finding in the minutes of a 17 June 1976 meeting:

> [M]eningitis appears to be more stable. Victor has run 60 patients through getting the correct answer for all of them. When a new patient was tried and the system did not get the right answer, Victor added new rules to make the patient run correctly – these additions were NOT detrimental to the system's performance on the other patients that already worked. In contrast, the 21 evaluation and pre-evaluation patients uncovered a number of bugs in the bacteremia rules. When rules were added or changed to fix those patients on which the system did not perform well, other patients that had worked previously now got the wrong answer. (This could be a result of the way the bugs were fixed rather than inherent instability of the bacteremia system.)[556]

Many of the challenges encountered during MYCIN's development stemmed from the patchwork nature of its creation. Shortliffe described this phenomenon in an email to his colleagues: 'I feel that comparisons of bacteremia and meningitis are suffering

[555] Clancey, '<CLANCEY>CF.TWOCUTOFF;1'.
[556] Carli Scott, '6/17/76 meeting summary', email printout, 18 Jun. 1976, ESP/10/'MYCIN-Current, 1975-82'.

somewhat from the simple fact that no one has ever sat down to try to make the entire bacteremia rule base consistent and coherent—those rules have developed over several years through an evolutionary process involving several rule authors'.[557] The bacteraemia and meningitis knowledge bases differed, Shortliffe argued, because of the social circumstances surrounding their creation.[558] The meningitis rule set was more amenable to revision because it was created by a small group of people in a relatively short period of time. Its makers had more of a global, synthetic understanding of the rule set and how it might respond to tweaks and edits. The bacteraemia rule set, by contrast, came into being over a period of multiple years—the product of countless authors and revisions. Shortliffe continued, 'Things could almost certainly be cleaned up if a single knowledgeable person devoted some time to revising the entire rule set. Meningitis, on the other hand, has been developed during a rather coherent short-term process during which the system itself has not undergone very many major changes'.[559]

TECHNOLOGIES OF PERSUASION

How could physicians know that the recommendations of a computer system like MYCIN were trustworthy? A rich answer to this question benefits from a brief foray into seventeenth-century England. Through the 1660s, the political philosopher Thomas Hobbes and the natural philosopher Robert Boyle were locked in a hot epistemological debate, one that centred on the validity of Robert Boyle's then-

[557] Shortliffe to Scott (Yu, Buchanan, David, van Melle, Clancey, Fagan, Axline, Wraith, and Aikins copied), 8 Jul. 1976.
[558] On the group dynamics around knowledge production and creation, see, e.g., Alma Steingart, 'A Group Theory of Group Theory: Collaborative Mathematics and the "Uninvention" of a 1000-Page Proof', *Social Studies of Science* 42, no. 2 (2012): 185-213.
[559] Shortliffe to Scott (Yu, Buchanan, David, van Melle, Clancey, Fagan, Axline, Wraith, and Aikins copied), 8 Jul. 1976.

unorthodox experiments with a newly invented device called the air-pump.[560] Where Boyle argued that the systematic manipulation of the natural world with machines like the air-pump represented the premier method of knowledge production, Hobbes, sceptical of inductive reasoning and concerned about the artificiality inherent in experimentation, instead looked to natural law for truths about nature. At stake in this dispute was nothing less than the legitimacy of experimentation as an accepted way of coming to know about the natural world.

As Steven Shapin and Simon Shaffer have famously shown, Boyle developed three 'technologies' designed to persuade others that the air-pump, and experimental methods more broadly, could truly reveal 'matters of fact': a material technology, involving the integrity of experimental components; a social technology, involving the norms by which competing claims are handled and a consensus is reached; and a literary technology, involving the communication of methods in a manner that allows readers to witness the experimental programme virtually, through written text.[561] Here I borrow from Shapin and Shaffer's rendering and frame the technologies built to convey the reasoning processes of computerised diagnostic systems to human physicians as literary technologies—as technologies of persuasion that served as a window for physicians into the unfamiliar reasoning processes of novel machines.[562]

---

[560] Steven Shapin and Simon Shaffer, *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life* (Princeton: Princeton University Press, 1985).
[561] Ibid., pp. 25-26.
[562] Shapin and Shaffer justify their branding of literary techniques as a 'technology' etymologically, invoking Plato's conception of *techne* which divided into two subcategories, 'one that consisted mainly of physical work and another that was closely associated with speech'. In the case of computer-aided diagnosis, these explanation capabilities were technologies both in this less familiar etymological sense and in a more familiar literal sense: not only a literary and social practice, the ability of a computer system to explain its reasoning processes was a product of technical feats in programming and engineering. Ibid., p. 25.

Designers of MYCIN, and other computerised systems, recognised early on the importance that machines be able to communicate, at least roughly, the steps by which they reached a given medical decision or diagnosis. The tight connection between the perceptibility and the credibility of an expert system's reasoning processes emerged as a theme of the 1975 Rutgers Artificial Intelligence (AI) in Medicine Workshop. Many attendees arrived collectively at the conclusion that only once a computerised system's reasoning processes were discernible to individuals without extensive training in computer science could physicians first evaluate, and subsequently choose to place their trust in, the medical recommendations of machines. The Rutgers computer scientist N.S. Sridharan voiced the strongest version of this conclusion, arguing that a computer-generated diagnosis, even if accurate, is functionally meaningless absent a mechanism by which to convey the logical processes underlying the diagnosis: 'It is not enough for the system to produce right answers', Sridharan declared. 'It has to be able to give reasons for those answers, in some sense explain its own processes. It has to be credible'.[563]

Many other participants chimed in, in agreement. The ophthalmologist Aran Safran, for example, commented on how the intelligibility of a computer's reasoning processes lay at the heart of 'the credibility of any medical system'.[564] Elaborating on this point, he explained, 'Given a data base and a set of numbers it is very important to be able to explain to a physician who is using them how these numbers were arrived at and their relative importance and how the system goes about reaching a decision or a hypothesis'. These same perspectives soon surfaced widely in the scientific and medical literature. That 'one of the primary requirements for user acceptance of a

[563] Amarel, Kulikowski, Sridharan, and Sridharan, *Proceedings of the First Annual AIM Workshop*, p. 30.
[564] Ibid., p. 37.

consultation program is an ability to explain decisions', that 'explanation capabilities in medical decision-support tools' are of paramount importance, that 'there has been an emphasis on giving [medical] advice systems an ability to explain their reasoning', and that 'producing explanations becomes an important task if the acceptability of the system is to be enhanced' reverberated through scientific journals and medical texts of all kinds.[565]

To achieve this goal of transparent reasoning, designers of different decision support systems created technologies that pursued a variety of strategies. What these technologies and strategies looked like depended in no small part on whether a given decision support system was built to supplement, imitate, or surpass human diagnostic reasoning. Ideas about the resemblance between a certain system's reasoning processes and human diagnostic reasoning informed how designers of these systems went about their efforts to make machine reasoning knowable to physicians. Some systems sought to recreate the reasoning practices exhibited by physicians generally, and sometimes even the more particular forms of reasoning exercised by specific practitioners. INTERNIST-I, for example, aimed to incorporate the medical expertise and imitate the logical processes of a single physician, the internist Jack D. Myers. Other systems sought to anchor diagnostic logic in something more fundamental, and less mercurial, than human medical practices: causal, pathophysiological models of disease. A notable exemplar of this approach, Rutgers University's CASNET programme, built its

[565] Edward H. Shortliffe, *Computer-Based Medical Consultations: MYCIN* (New York, 1976), p. 195; Edward H. Shortliffe, 'Computer Programs to Support Clinical Decision Making', *Journal of the American Medical Association* 258, no. 1 (1987): 61-66, p. 64; Edward H. Shortliffe, 'Medical Expert Systems-Knowledge Tools for Physicians', *Western Journal of Medicine* 145, no. 6 (1986): 830-839, pp. 833-834; Sholom Weiss, Casimir A. Kulikowski, Saul Amarel, and Aran Safir, 'A Model-Based Method for Computer-Aided Medical Decision-Making', *Artificial Intelligence* 11 (1978): 145-172, p. 146.

diagnostic logic around foundational knowledge about 'the pathogenesis and mechanisms' of optic nerve diseases, represented 'in terms of cause and effect relationships between pathophysiological states'.[566]

The makers of MYCIN occasionally emphasised that their programme was more ends-oriented than means-oriented. As Shortliffe said of the MYCIN programme, 'we are not necessarily trying to make the program perform the way a clinician does'.[567] Rather, he continued, 'we want the program to derive the right advice and whatever way we can come up with to do that is all right'. Yet, even if the MYCIN team was not set on creating a system that replicated *exactly* human forms of medical reasoning, they remained intensely interested and invested in the parallels between MYCIN's reasoning processes and those of physicians.

In particular, MYCIN's rule-based approach, with medical knowledge encoded into discrete information packets that took the form of production rules, was decided on under Shortliffe's assumption that this method promised 'a straightforward mechanism for explaining decisions'.[568]

As deeply as the principle of intelligibility was built into MYCIN's knowledge base and inference engine, designers still believed that, in order to make the computer system convincing, it required a manner by which to accessibly explain its decisions. Even when first envisaging MYCIN through the early 1970s, Shortliffe appreciated that diagnostic accuracy alone would not be enough to secure acceptance among medical practitioners—that 'other issues such as mode of interaction, transparency of the "reasoning" behind a system's advice, and integration of a system into the physician's

---

[566] Weiss, Kulikowski, Amarel, and Safir, 'A Model-Based Method for Computer-Aided Medical Decision-Making', p. 147
[567] Amarel, Kulikowski, Sridharan, and Sridharan, *Proceedings of the First Annual AIM Workshop*, p. 23.
[568] Shortliffe and Buchanan, 'A Model of Inexact Reasoning in Medicine', p. 357.

daily routine also seemed to be crucial determinants of acceptance'.[569] Of these 'other issues' feeding into the system's clinical viability, Shortliffe and his collaborators quickly found themselves 'particularly interested in explanation and…drawn to the techniques of artificial intelligence as potential methods for developing consultation programs that could justify their advice'.[570] The clarity of discrete medical rules did not go far enough. Though Shortliffe discovered that MYCIN's production rules, 'if translated into English, could provide explanations that were usually understandable', the mere 'regurgitation of rules', he concluded, 'is not always satisfactory'.[571] The *manner* in which these rules were presented mattered as far as the system's intelligibility and credibility were concerned. This problem of explanation, the MYCIN team believed, had a technological solution.

It was the observation about the elucidative limits of production rules, particularly when these rules were presented with little context or justification, that motivated the Stanford team to incorporate an explanation mechanism into the larger MYCIN system. The defining feature of MYCIN's explanation facility, dubbed unromantically 'Subprogram 2', was its interactive character. Rather than generically producing a summary of its logical processes after each consultation session, MYCIN instead revealed its underlying logics only in response to specific, targeted questions posed by physicians. As the consultation progressed from the decision support phase to the explanation phase, from Subprogram 1 to Subprogram 2, the responsibility for actively directing the interaction therefore shifted from computer to physician. Shortliffe described this role reversal in his 1976 book on the pioneering expert system:

---

[569] Edward H. Shortliffe, 'Clinical Consultation Systems: Designing for the Physician as Computer User', *Proceeding of the Annual Symposium on Computer Applications in Medical Care* (1981): 235-236, p. 235.
[570] Ibid., p. 235.
[571] Ibid., p. 235.

'Unlike the Consultation System (Subprogram 1) in which MYCIN takes the initiative, asking questions and waiting for the physician to respond, the Explanation System expects the user to guide the interaction'.[572] To portray the explanation phase of the MYCIN consultation as entirely physician-lead, however, is to overstate the physician's degree of autonomy in his engagement with the computer system: the kinds of questions that MYCIN was capable of fielding were limited, and even the subset of questions that *could* be asked also had to conform to various syntactic and structural constraints.

In order to make MYCIN's explanations and justifications intelligible to physicians with little training in programming or computer science, Shortliffe and his colleagues first needed a way of translating the system's production rules, internally coded in the LISP programming language, into natural English. Because of the structured nature of LISP, the MYCIN team could generate a meaningful translation by associating LISP functions with translation patterns.[573] As a result, what was coded internally as

```
PREMISE:  ($AND  (SAME CNTXT GRAM GRAMNEG)
                 (SAME CNTXT MORPH ROD)
                 (SAME CNTXT AIR ANAEROBIC))
ACTION:   (CONCLUDE CNTXT INDENT BACTEROIDES TALLY .6)[574]
```

would appear to the user as

```
IF:    1) THE STAIN OF THE ORGANISM IS GRAMNEG, AND
       2) THE MORPHOLOGY OF THE ORGANISM IS ROD, AND
       3) THE AEROBICITY OF THE ORGANISM IS ANAEROBIC
THEN:      THERE IS SUGGESTIVE EVIDENCE (.6) THAT THE
           IDENTITY OF THE ORGANISM IS BATERIODES[575]
```

---

[572] Shortliffe, *Computer-Based Medical Consultations: MYCIN*, p. 197.
[573] Ibid., pp. 78-132.
[574] Edward H. Shortliffe, Randall Davis, Stanton G. Axline, Bruce G. Buchanan, C. Cordell Green, and Stanley N. Cohen, 'Computer-Based Consultations in Clinical Therapeutics: Explanation and Tule Acquisition Capabilities of the MYCIN System', *Computers in Biomedical Research* 8, no. 4 (1975): 303-320, p. 306.
[575] Ibid., p. 305

This method of translating LISP into English represented a reverse engineering of the way in which medical knowledge was originally coded into a novel digital medium. As the computer scientist William van Melle described it in 1984, 'this process involves having the medical expert supply a piece of medical knowledge in English, which a system programmer converts into the intended LISP rule'.[576]

The Stanford team then turned their attention to the tasks of functionalising MYCIN's explanation capabilities and integrating them into the broader operation of the decision support system. The creators of MYCIN designed the system's explanation function such that physicians could tap these capabilities at two different points, either *during* the consultation session or *after* advice is given. To examine MYCIN's reasoning chain during a consultation, physicians had a set of commands at their disposal, primarily intended to help the user understand what the programme was after in asking particular questions throughout the consultation. By typing, for example, 'WHY' into the MYCIN system, physicians could elicit an explanation of a particular question's 'purpose'—that is, the question's relationship to the subgoal that MYCIN was trying to reach. Understanding the purpose of a given question, however, did not guarantee full clarity: even if a physician grasped how a question served MYCIN in working toward a particular subgoal, he may not fully comprehend the programme's reason for seeking that particular subgoal in the first place. To further enquire about the underlying subgoals, and their connection to the system's main goal, physicians could type 'WHY' into the system a second time (Figure 26). Physicians could also trigger

[576] William van Melle, 'The Structure of the MYCIN System', in Bruce Buchanan and Edward Shortliffe (eds.), *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (Reading, MA: Addison-Wesley Publishing Company, 1984), 67-77, p. 76.

explanations from MYCIN by using the HOW command. Unlike the WHY command, which elucidated MYCIN's active reasoning chain, the HOW command allowed users to explore other possible reasoning chains—to investigate, for instance, the different ways a given subgoal could hypothetically be achieved.

MYCIN's designers assumed that, even with explanatory support throughout the consultation phase, physicians might still be mystified by the system's eventual medical recommendation. Often MYCIN drew inferences without having requested any recent information from the physician, something that could happen, the MYCIN team explained, 'either because the preconditions of a rule were satisfied by deductions from previous information (using other decision rules) or because the needed data had already been requested in the context of a different rule'.[577] As a result, if physicians understood the meaning and purpose of each question in isolation, by no means did this assure that they would understand the system's final recommendation. The explanation capabilities initiated after the conclusion of the consultation session was developed to address this concern. MYCIN's post-consultation explanation capability took the form of a question-answering (QA) session, during which the physician could pose queries concerning the current line of reasoning, MYCIN's production rules more generally, MYCIN's reasons for asking a previous question, and the accuracy of the physician's own production rules (e.g., 'Is chloramphenicol okay for Salmonella infections?').

A common assumption among designers and users of decision support systems was that rule-based systems like MYCIN held a categorical advantage over purely statistical methods in explaining their reasoning processes and medical knowledge to physicians in a natural and efficient manner. The Swedish linguist Hans Karlgren

[577] Shortliffe, Davis, Axline, Buchanan, Green, and Cohen, 'Computer-Based Consultations in Clinical Therapeutics, p. 314.

articulated this view in a review of the 1977 book *Computational Linguistics in Medicine*, writing in reference to MYCIN, 'The rule approach has its limitations but provides an impressive explanation capability, without addressing all the complex problems of computational linguistics'.[578] Those who worked directly on rule-based approaches to computerised decision support in medicine similarly commented on the intrinsic intelligibility of this approach, especially when contrasted with earlier, statistical approaches. The CASNET team, for example, underscored the conceptual distinction between their rule-based diagnostic methods and alternative statistical methods, writing 'AI techniques attempt to capture decision making rules explicitly, while statistical methods may extract them implicitly from accumulated sample experience'.[579] The former approach, involving 'a more structured representation of the diagnostics and therapy selection problems', was far better suited to clear and effective explanation—these systems, the team reasoned, 'are more likely to be accepted because they are expressed in a decision-making context familiar to the clinician'.[580]

These built-in advantages of rule-based methods of computerised diagnosis notwithstanding, MYCIN's explanation capabilities were not entirely without limitations, nor did these capabilities come without certain costs. Chief among such costs involved the fundamental trade-off between intelligibility and efficiency: improvements in the system's explanation competencies could come at the expense of its speed. Shortliffe broached this dilemma in a paper delivered at the 1977 Annual

---

[578] Hans Karlgren, Review of *Computational Linguistics in Medicine*, in Werner Schneider and Anna-Lena Sågvall Hein (eds.), *Journal of Pragmatics* 4, no. 2 (1980), p. 183. G. Anthony Gorry was among the first to make this observation. See G. Anthony Gorry, 'Computer-Assisted Clinical Decision Making', *Methods of Information in Medicine* 12 (1973): 45-51.
[579] Weiss, Kulikowski, Amarel, and Safir, 'A Model-Based Method for Computer-Aided Medical Decision-Making', p. 162.
[580] Ibid., p. 162.

Meeting of the Society for Computer Medicine: 'The efficiency advantages of approaches relying more heavily on statistical theory must be weighed against the importance of the natural mechanism for explanation and knowledge modification that may be achieved through MYCIN's representation of knowledge in the form of production rules'.[581] Achieving the appropriate balance between these competing values proved evasive, as intelligibility and efficiency, though clashing aims, both promoted the computer system's credibility in separate yet equally essential ways. The question of whether to pursue a system that was more accurate or more intelligible had no technological solution; this question occupied the realm of values and morals.

A number of technical and intellectual barriers compounded these concerns over diagnostic efficiency. Among the largest barrier was the nascent state of natural language processing, a subfield of computer science concerned with the production and comprehension of natural human language. While Subprogram 2 could translate the production rules, internally coded in LISP, into English, these translations, the MYCIN team conceded, were 'at times stilted because natural language processing is not a focus' of their research.[582] Not only an issue of communicating to physicians, the underdeveloped state of natural language processing bred issues around MYCIN's ability to understand physicians as well; in the MYCIN team's words, this was 'the problem of computer understanding of natural language'.[583] Attuned to the close relationship between natural language processing and the credibility of expert systems in medicine, designers of the systems agreed that improvements in latter would come

[581] Edward Shortliffe, 'MYCIN: A Knowledge-Based Computer Program Applied to Infectious Diseases', *Proceeding of the Annual Symposium on Computer Applications in Medical Care* (1977): 66-69, p. 67.
[582] Shortliffe, Davis, Axline, Buchanan, Green, and Cohen, 'Computer-Based Consultations in Clinical Therapeutics', p. 305.
[583] Ibid., p. 315.

hand-in-hand with improvements in the former: 'The problem we are all experiencing in trying to explain our own systems and understanding others', N.S. Sridharan stated, 'could be alleviated also if the knowledge were represented in a more readable, natural form. It would make it much easier to get down to the concrete stuff of the system and follow its reasoning processes'.[584]

Although designers concluded that advancements in natural language processing would go a long way toward producing a system that could interact with physicians seamlessly, clearly, and compellingly, they knew that such developments could only go so far: no matter how polished a system's explanation capabilities, to heed or not to heed the system's advice was a choice that ultimately rested with the physician. This observation arose as a topic of conversation at the inaugural AI in Medicine Workshop, with Shortliffe noting, 'The expert may ask us how the program reached the conclusion and we can list for him the six rules. And each one of those rules may look just great to him but he simply cannot accept the conclusion'.[585]

- CHAPTER VI -

[584] Amarel, Kulikowski, Sridharan, and Sridharan, *Proceedings of the First Annual AIM Workshop*, p. 30.
[585] Ibid., p. 37.

‘Hidden in the Code’

MYCIN’s ability to explain itself was about more than transparency and trust. It was also about authority. By opening up MYCIN’s reasoning process to scrutiny, Shortliffe and his colleagues were affirming the physician’s role as the ultimate decision maker. To lay bare the computer’s reasoning before the physician was as performative an act as it was a practical one. It communicated to the physician that, for all the complex medical decisions that the computer made, it was up to him to decide whether or not any given decision would stand. In publications and user manuals, the MYCIN team frequently hinted at this connection between explanation and authority. Advice that could be understood, they often noted, was advice that could be reviewed and, ultimately, rejected. A 1976 version of the MYCIN User’s Manual, for example, noted that the explanation feature not only ‘permit[ted] the program to justify its conclusions’ but also ‘allow[ed] the user to reject the advice if he feels that some step in the reasoning process has been unsound’.[586] Once Shortliffe joined Stanford’s faculty and began overseeing his own research group, some of his students carried forward this line of thinking. One of them, Greg Cooper, pursued, as a major thread of his PhD work, research into computerised explanation systems, the motivation behind which, he noted in a message to Shortliffe, ‘is the deeply hel[d] belief that physician’s want active control of the diagnostic process…’.[587]

That MYCIN’s explanation capabilities induced physician interaction and engagement was no accident. Interaction, MYCIN’s designers recognised, nurtured a sense of participation and ownership, ‘meet[ing] many physicians’ desires’ to feel as

---

[586] ‘MYCIN USER’S GUIDE’, 9 Mar. 1976, ESP/10/‘MYCIN-Current, 1975-82’.
[587] Gregory Cooper to Edward Shortliffe, Bruce Buchanan, and csd.lenat@SCORE, email printout, 16 Mar. 1982, EHS/1/‘Cooper, Gregory, 1979-1996’.

though they are 'formulat[ing] [treatment] plans on their own'.[588] Subprogram 2, then, did more than allow physicians to virtually witness machine inferences; the technology fostered a feeling of involvement in and ownership over these novel styles of reasoning. The interactive dimension of Subprogram 2 advanced the explanation system's broader aim of 'clarify[ing] and emphasiz[ing] the physician's role as the ultimate decision maker in patient management' and of 'acknowledg[ing] that the users are ultimately making the decision and are using the computer program as an adjunct'.[589] A principal function of MYCIN's explanation capability—its literary technology—was the inscription of hierarchy.

Because explanation signalled deference, MYCIN's creators saw the explanation capability as valuable even if it went unused. In 1974, Homer Warner, who was serving as the inaugural editor of the new journal *Computers and Biomedical Research*, wrote to Shortliffe about the necessary revisions to a manuscript that the MYCIN team had submitted to the journal. Of particular interest to Warner and his reviewers was MYCIN's explanation capability. 'What is the basis', Warner probed, 'for the statement that Doctors will not use advice given by a computer unless a thorough explanation of the basis for this advice is also provided? For instance, how often do they use the explanation program? What is the total number of cases on which the consultation program has been applied to date?'[590] Shortliffe responded to Warner's inquiry candidly. He confessed that preliminary observations of the system in experimental settings seemed to classify MYCIN's explanation capability as marginal—at least as measured by the yardstick of direct utilisation: 'Based on

[588] Shortliffe, 'Computers to Support Clinical Decision Making', p. 63.
[589] Shortliffe, 'Medical Expert Systems-Knowledge Tools for Physicians', p. 833; Shortliffe, 'Computers to Support Clinical Decision Making', p. 64.
[590] Homer Warner to Shortliffe, 6 Dec. 1974, ESP/12/'CMPB (Computer Methods and Programs in Biomed.), 1974-94'.

observation of individuals who have experimented with the system to date, we do not expect that its ability to explain will be used with regularity by most users'.[591] But the explanation capability, Shortliffe continued, remained vitally important because it dictated the ontology of MYCIN; it shaped what MYCIN was. Sans an ability to explain itself, MYCIN risked becoming an inflexible, top-down dictator of medical advice. Cum such an ability, however, MYCIN became just another useful tool over which the physician retained control. As Shortliffe explained to Warner, 'The capability is important, however, because it makes the program function as a tool to the physician rather than as a dogmatic advice giver'.[592] Whether MYCIN *could* explain itself was, for the MYCIN team, more essential than whether MYCIN, in any given clinical encounter, *did* explain itself.

Despite all that was said and done to achieve deference, in reality MYCIN's relationship to the physician-user and his traditional authority was far more muddled. In a retrospective appraisal of the MYCIN project, programmer William Clancey argued that the system's mode of consultation was inescapably hierarchical, subordinating the physician-user's knowledge to that of MYCIN's architects. He wrote, 'We viewed all interactions between people and program in terms of "transfer": Experts had knowledge stored in their heads. The knowledge base was an objective inventory of expert knowledge. Users lacked knowledge. The role of consultation … was to transfer knowledge between experts, users, and students'.[593] For Clancey, such a unidirectional model of transfer had an inherently tiered, even religious, character. 'Knowledge engineers', he ventured, 'are like priests; they receive "The Word" from

[591] Shortliffe to Warner, 31 Dec. 1974, ESP/12/'CMPB (Computer Methods and Programs in Biomed.), 1974-94'.
[592] Ibid.
[593] William J. Clancey, 'Notes on "Epistemology of a Rule-Based Expert System"', *Artificial Intelligence* 59, no. 1-2 (1993): 197-204.

experts above, add nothing to the content, but codify it accurately into written rules, and pass it down to ordinary folks as commandments to live by. If you are not an expert, you learn by being told'.[594] Such a view of knowledge transfer, Clancey maintained, reinforced an impoverished view of medical knowledge—it 'lack[ed] any sense that the knowledge base could belong to a community of practitioners, the users and experts alike, and might be developed and maintained by them'.[595]

To a certain extent Clancey was selling himself and his Stanford colleagues short. The team did in fact see physician-users as potential sources of additional medical knowledge and even designed capabilities for physicians to edit and expand the existing body of rules. In perceiving users as potential ancillary fonts of knowledge, MYCIN members also recognised the limitations to their own knowledge. They held informed, if not sophisticated, views about the nature of medical knowledge; they understood that MYCIN's rule-based conception of knowledge was imperfect and shallow—that not every aspect of medical expertise was explicit, transferrable, and rule-like. Indeed, their thinking on such matters was informed by contemporary research in the history and philosophy of science. One member observed, citing Polanyi's then-recent work on the tacit and personal dimensions of knowledge, that not all knowledge in medicine was likely to be amenable to MYCIN-like formalisms: 'We might argue from sources like Kuhn and Polanyi that the 'tacit' component of knowledge is gained only from direct experience with the task …'.[596] (This team member nevertheless maintained that surely 'sometimes human thinking IS rule-like'.)

---

[594] Ibid.
[595] Ibid.
[596] 'The Student Model and the Form of the Knowledge Base', unpublished notes, n.d., EHS/2/'Clancey, William (1), 1978-81'. No author is listed for this document, although it was surely produced by a member of the MYCIN team. I believe it is William Clancey on the basis of its content, style, and location within Shortliffe's papers.

The MYCIN team, then, did not espouse a wholly naïve and unchecked view about the nature of medical knowledge. Their attempts to explicate medical knowledge derived from a pragmatic impulse to build, under conditions of limited time and funding, a system that worked, as well as their view that medical knowledge, even if 'tacit' in some regards, was at least partially formalizable.

Though an oversimplification, Clancey's retrospective critique did get some things right. For all that the MYCIN team said publicly about deference, behind closed doors the creators of MYCIN sometimes did portray the physician-users as, to use Clancey's words, 'ordinary folks'. At times they presumed clinicians to be ignorant and unskilled. A significant internal consideration among the MYCIN team concerned the possibility of designing the system so as to minimise the physician-user's clinical judgement. The kinds of information that MYCIN requested from its physician-user arrayed along a spectrum of clinical discretion. Certain kinds of patient information required next to no clinical judgement to obtain. Data from the microbiological laboratory, for instance, could simply be transferred from the lab report into MYCIN. But other kinds of information about the patient derived from the clinician's own judgement—his personal evaluation of the patient.

Questions nagged at some members of the MYCIN team about the appropriate degree of trust to place in the physician's ability to accurately obtain these more judgemental forms of patient information. MYCIN, after all, was predicated on the assumption that its user lacked expertise in the area of infectious disease diagnosis and treatment. If the physician-user needed assistance in the realm of infectious disease, could he really be trusted to discern, evaluate, and report clinic findings related to this area of medicine? Questions of this kind were, Shortliffe recounted in a 1976 email, the topic of 'several sessions over the years when we have had to discuss the issue of how

much judgmental information to request from the presumed non-experts using our system'.[597] Much thought went into, as another MYCIN affiliate put it, 'avoid[ing] the problem of user variability and fallibility'.[598] The rationale for minimising reliance on clinical judgement was best summarised in the minutes of a 1976 MYCIN meeting: 'The closer the system is tied to objective information, and the fewer judgmental Q's it asks, the more reliable it will be. The MD's who most need help from MYCIN are the least likely to give sound answers to the judgmental Q's'.[599]

Considerations about the proper balance between judgemental and non-judgemental forms of medical knowledge, often discussed in the abstract, gained a more concrete figure as MYCIN affiliates turned their attention to one particular patient finding used by MYCIN: a blood infection's portal of entry. The portal through which a given organism entered the human body—whether broken skin, the respiratory tract, mucous membranes, the genitourinary tract, or catheters—could, on occasion, help clinicians (and MYCIN) identify the infectious organism. Yet, at least one MYCIN affiliate worried that the portal concept demanded too much of the non-expert physician: '[T]he worst flaw is that its [the portal concept's] use requires clinical judgment and knowledge on the part of the user, which he may or may not have. If the user is sophisticated, the selection of the portal may be "correct" and Mycin has benefits from his knowledge. Likewise, if the user selects an inappropriate portal, then Mycin and the patient will suffer the consequences of his error. The consultant (Mycin)

---

[597] Shortliffe to Scott (Yu, Buchanan, David, van Melle, Clancey, Fagan, Axline, Wraith, and Aikins copied), email printout, 8 Jul. 1976, ESP/10/'MYCIN – Current, 1975-82'.
[598] 'Critique of the Portal Concept in Mycin', unpublished document, n.d. ESP/10/'MYCIN-Current, 1975-82'. The member of the MYCIN team who authored this document is unknown. It is possibly Victor Yu, although this is a low-confidence speculation.
[599] Buchanan, 'Summary of 11/17/76 meeting', meeting minutes, printout of email to MYCIN team, 30 Nov. 1976, ESP/10/'MYCIN-Current, 1975-82'.

then becomes dependent on the user'.[600] The concern here was that MYCIN's reliance upon judgemental information left its performance up to chance: There, but for the grace of God, MYCIN goes.

For Shortliffe's part, he agreed that ideally MYCIN ought to rely upon as little judgemental information as possible: '[I]n an ideal world we would like to ask the user only for LABDATA'—that is, non-judgemental information that came from the microbiological laboratory. But given that MYCIN would be operating under non-ideal conditions, he was more willing than some of his colleagues to cut the physician-users some slack. Shortliffe hastened to remind colleagues that the philosophy behind MYCIN—one of deference to physicians—aligned with the system's dependence, at least partially, on the user's judgement, skill, and competence: 'I personally feel that we should give the nonexperts at least SOME credit for being able to make useful observations. It is, after all, consistent with our initial philosophy of MYCIN that we not try to replace entirely the physician's reasoning powers but rather to augment them'.[601]

These discussions about clinical skill and discretion illustrate how considerations of trust and trustworthiness applied to MYCIN's users as much as they applied to MYCIN itself. When the team focused on MYCIN's clinical acceptance, their primary concern was building the physician's trust in the system. For traditional physicians to accept a foreign technology like MYCIN, its recommendations (and the chains of reasoning behind them) had to be trustworthy—understandable by humans and open to verification, revision, and ultimately dismissal. But when the team shifted their attention to the issue of accuracy rather than that of acceptance, their concern

[600] 'Critique of the Portal Concept in Mycin'.
[601] Shortliffe to Scott (Yu, Buchanan, David, van Melle, Clancey, Fagan, Axline, Wraith, and Aikins copied), 8 Jul. 1976.

expanded to include the competence and trustworthiness not just of the system but of its human user as well. Was MYCIN's imagined user clinically skilled enough to accurately obtain certain kinds of patient information? In attempting to build a system that they believed would be both deferential and accurate, Shortliffe and his colleagues often found themselves straddling a fine line. At the same time that members of the MYCIN team were holding the physician up as the ultimate decision maker, they were also privately questioning his clinical ability.

MYCIN IN THE ANTIBIOTIC ERA

The relationship between MYCIN and clinical authority must be understood against the backdrop of much larger tensions and transformations around antibiotics and their irrational use. While MYCIN was often thought of as a form of basic research in computer science and artificial intelligence, a second aim of the enterprise was applied and clinical: to engender a rational approach to antibiotic prescribing. In its clinical objectives, MYCIN belongs to a larger and older reform effort, one whose origins date back at least two decades before MYCIN's creation. As Scott Podolsky has shown, antibiotic reform efforts first coalesced in the 1950s—as much the result of growing anxieties about antibiotic misuse as the emergence of the infectious disease specialist as a new category of medical expert.[602] In these early years, antibiotic reformers directed their energies to a class of drugs that are now largely forgotten: fixed-dose combination antibiotics. Pharmaceutical companies promoted this class of drugs, combinations of multiple antibiotics into a single pill, on the basis that the whole

[602] Scott Podolsky, *The Antibiotic Era: Reform, Resistance, and the Pursuit of a Rational Therapeutics* (Baltimore: Johns Hopkins University Press, 2015), 10-42. This section's discussion of mid-century antibiotic reform is very much beholden to Podolsky's book.

was greater than the sum of its parts. But reformers remained sceptical. Infectious disease experts, led by Harvard's Maxwell Finland among others, pointed out that assertions about the drugs' worth rested on first-person testimonials rather than controlled clinical trials.[603] These critiques proved persuasive: reformers would see the fruits of their labour in 1962, with the enactment of the Kefauver-Harris amendments. The amendments granted the FDA the power to remove from the market drugs, like fixed-dose combination antibiotics, that had not been proven (via 'adequate and well-controlled investigations') to be efficacious.[604]

The passage of the Kefauver-Harris amendments freed antibiotic reformers to focus on a different aspect of medical practice that needed reforming: physicians' prescribing practices. By the 1970s, the attention lavished on the irrational prescribing practices of many physicians reached something of a crescendo, as reformers worked to underscore the problems associated with antibiotic overuse in medical, governmental, and popular fora. Beginning in 1967, Wisconsin Senator Gaylord Nelson initiated widely publicised hearings on the pharmaceutical industry—in many ways a follow-up to the Kefauver hearings that had transpired nearly a decade earlier. Initially, the hearings probed a broad set of issues, ranging from the continued challenges of pharmaceutical competition and drug pricing to newer concerns about the safety of oral contraceptives.[605] But in 1972, the same year that Shortliffe began work on MYCIN,

---

[603] Ibid., 43-72.

[604] Podolsky, *The Antibiotic Era,* 73-111. For a brief but rich historical discussion of the Kefauver-Harris amendments, see Jeremy Greene and Scott Podolsky, 'Reform, Regulation, and Pharmaceuticals — The Kefauver–Harris Amendments at 50', *NEJM* 367, no. 16 (2012): 1481-1483. See also Daniel P. Carpenter, *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA* (Princeton: Princeton University Press, 2010), pp. 228-297.

[605] On the Nelson hearings, see Dominique Tobbell, *Pills, Power, and Policy: The Struggle for Drug Reform in Cold War American and Its Consequences* (Berkeley: University of California Press, 2011), 129-130, 163-192; Jeremy Greene, *Generic: The*

the hearings intensified their focus on the misuse and overuse of antibiotics. Additional high-profile Congressional hearings into antibiotic overuse, led by Massachusetts Senator Edward Kennedy, continued over the following two years.[606] In their appearances before these hearings, and in their broader campaign to change prescribing behaviours, antibiotic reformers frequently resorted to the rhetoric of 'rationality'.[607] Theirs was an effort to banish rampant 'irrational' prescribing and to cultivate in its place a 'rational' approach to antibiotic use. Although antibiotic resistance would later come to serve as the chief impetus driving such reform efforts, in the early 1970s reformers were motivated by a host of concerns beyond antibiotic resistance alone: reformers argued that irrational prescribing produced superinfections, adverse side effects, costs, and even a broader 'dumbing down' of the medical profession.

The topic of antibiotic misuse also received heightened attention in less public settings during this period. In medical journals, many researchers turned a critical—and empirical—eye to physicians' prescribing behaviours, producing studies that quantified the irrationality and regional variation in antibiotic use.[608] One influential 1972 study of a 500-bed community hospital found 66 percent of the hospital's antibiotic prescriptions to be 'irrational' as judged by a panel of expert physicians and pharmacists.[609] Only 13 percent of therapeutic decisions in the realm of antimicrobials

---

Unbranding of Modern Medicine (Baltimore: Johns Hopkins University Press, 2014), pp. 53-55; Carpenter, *Reputation and Power*, 342-345.

[606] Podolsky, *The Antibiotic Era*, p. 136.

[607] Ibid., pp. 2-4, 38.

[608] See, e.g., W.E. Scheckler and J.V. Bennett, 'Antibiotic Usage in Seven Community Hospitals', *JAMA* 213 (1970): 264-267; K.E. Reztak and R.B. Williams, 'A Review of Antibiotic Therapy in Patients with Systematic Infections', *American Journal of Hospital Pharmacy* 29 (1972): 935-941; C.M. Kunin, T. Tupasi, and W.A. Craig, 'Use of Antibiotics: A Brief Exposition of the Problem and Some Tentative Solutions', *Annals of Internal Medicine* 79 (1973): 555-560.

[609] A.W. Roberts and J.A. Visconti, 'The Rational and Irrational Use of Systematic Antimicrobial Drugs', *American Journal of Hospital Pharmacy* 29 (1972): 828-834.

met the standard of 'rationality'. Reporting similar findings, another study, published in a 1974 issue of *JAMA*, concluded that '[t]he gap between the actual antibiotic prescribing practices and the ideal practices recommended by infectious disease specialists appears to be widening'.[610] Different regulatory bodies started to mobilise—and various guidelines, task forces, and committees started to materialise—in response to rising levels of concern about antibiotic misuse. The MYCIN team followed these studies closely and cited them as evidence of a need for a computerised consultation system such as theirs. The weight of empirical evidence, the MYCIN team observed, 'reflects the need for improved therapy selection in patients requiring therapy. This', they continued, 'is precisely the decision task with which MYCIN is designed to assist'.[611]

Such developments point to an important evolution in the nature of therapeutic reform efforts over the second half of the twentieth century, as Podolsky has demonstrated. In the 1950s and 1960s, therapeutic reformers primarily targeted irrational drugs, arguing that fixed-dose combination drugs were not just inefficacious but potentially dangerous. Having successfully worked to remove these drugs from the market, therapeutic reformers shifted their focus. Beginning in the late 1960s, reformers increasingly worried not about irrational drugs but about the irrational use of drugs that, in some contexts, were entirely warranted and efficacious.[612] The main target of reform, in other words, changed from the faulty *drugs* in the physician's armamentarium to the faulty *behaviours* of physicians themselves. As a decision

---

[610] Henry Simmons and Paul Stolley, 'This Is Medical Progress? Trends and Consequences of Antibiotic Use in the United States', *JAMA* 227 (1974): 1023-1028, p. 1027. Increasing recognition of unwarranted and widespread variation in medical practices and outcomes during the 1970s and 1980s extended beyond the realm of antibiotics. See Jones, *Broken Hearts,* pp. 215-218.
[611] Shortliffe, *Computer-Based Medical Consultations: MYCIN*, p. 43
[612] Ibid., p. 112.

support tool that aimed to intervene in physicians' prescribing practices, MYCIN aligned with this new reform emphasis. In fact, the makers of MYCIN sometimes adopted the very vocabulary of therapeutic reformers, describing their system as a way of shoring up 'rational', and reining in 'irrational', prescribing behaviours.

MYCIN wasn't the only effort during this period to bring new computational tools to the challenge of antibiotic decision making. At Harvard's Peter Bent Brigham Hospital, a team led by Thomas O'Brien used a computer programme to analyse and plot the results of antimicrobial susceptibility tests.[613] Beginning in March 1966, the group started encoding the results of all susceptibility tests conducted at the hospital's microbiology laboratory onto 'computer cards'. By 1969, they had amassed over 20,000 cards. This data was then analysed by computer and rendered into separate plots depicting the susceptibility of different microbes to different antibiotics.[614] The team then arranged these plots into a larger grid, with microbes on the x-axis and antibiotics on the y-axis (Figure 27). In doing so, O'Brien and his colleagues created a tool that allowed physicians to glean (by scanning horizontally) which antibiotics a given organism might be susceptible to and (by scanning vertically) which organisms a given drug might be effective for. Ultimately the team hoped that these computer-generated plots would assist physicians in quickly making judgements about antibiotic selection on the basis of the hospital's statistical patterns of susceptibility. Other, analogue technologies also emerged to assist physicians prescribe appropriate antibiotics—and promote antibiotic stewardship. Notable among these was Jay P. Sanford's guide to antimicrobial therapy. Originally created as a 1969 handout for a medical grand rounds

[613] Thomas F. O'Brien, Ralph L. Kent, and Antone A. Medeiros, 'Computer-Generated Plots of Results of Antimicrobial-Susceptibility Tests', *JAMA* 210, no. 1 (1969): 84-92.
[614] The zone of inhibition is the result of an antibiotic effectively killing or stopping the growth of the bacteria on the agar plate.

presentation, the Sanford guide evolved into a formally published booklet that physicians carried in their white coat pockets—often with the same ubiquity with which they wore stethoscopes around their necks.

The linguistic and ideological overlap between antibiotic reform and the MYCIN project becomes particularly apparent in a 1977 exchange between the HEW's National Center for Health Services Research and Stanley Cohen. In late March of that year, Richard DuBois of the HEW delivered news to Cohen that the National Center's study section recommended only a partial renewal of funding for the MYCIN project. The study section's decision reflected their concerns about the system's clinical utility. Specifically, reviewers seemed to believe that the selection of infectious disease as MYCIN's problem domain placed inherent limitations on MYCIN's potential usefulness in medical practice:

> While regarding MYCIN as an excellent model of clinical decision-making for educational purposes, reviewers questioned its utility in medical practice. One factor contributing to reviewers' reservations is that the choice of antibiotics in practice is far more constrained than the choices built into MYCIN. Institutional norms and policies often dictate which antibiotics to use, and the necessity of protecting the patient against unidentified organisms often serves as justification for using a broad-spectrum antibiotic.[615]

This conclusion did not sit well with Cohen. Channelling the energies, insights, and even anger of the leading antibiotic reformers, Cohen shot back in a heated letter to the National Center's director, 'I am astounded that the Center would officially take the position, as it seems to have done, that there is not really much need for developing a rational approach to antimicrobial therapy—since we can simply give broad-spectrum antibiotics'.[616] When Cohen expanded upon these critiques in the same letter, he

[615] Richard DuBois to Stanley N. Cohen, 30 Mar. 1977, ESP/10/'MYCIN History, 1972-78'.
[616] Cohen to Geral Rosenthal, 8 May 1977, ESP/10/'MYCIN History, 1972-78'.

adopted the language of antibiotic reformers, contrasting MYCIN's 'rational approach' with the reviewers' 'irrational' position:

> This claim is complete rubbish; if selection of a broad-spectrum antibiotic to treat various infectious diseases is all that is necessary, there would be no need for diagnostic laboratories to culture organisms from diseased patients, and no need for infectious disease consultations by either human experts or computers. A major sub-specialty of Internal Medicine, Infectious Diseases, could simply disappear if the problem were as simple as some observers seem to believe. The view expressed is a prime example of the misinformation and ignorance that leads every year to the waste of billions of dollars of health care funds for irrational and unnecessary antibiotic use.[617]

According to Cohen, then, the study section got things completely backwards: the fact that 'institutional norms and policies often dictate' antibiotic use was proof not of MYCIN's futility but of its clinical relevance and urgency. Physicians' irrational prescribing practices, driven by 'misinformation and ignorance', were in need of correction.

Unsurprisingly, many physicians did not respond lightly to the suggestion that they were major contributors to the problem of antibiotic resistance. Indeed, many denied that there was a problem at all. 'Yes, we are 'over-prescribing' antibiotics, and it is good that we are', opined one clinician in 1974. 'My God, let's not go back to the "good old days"!'[618] Of course, such sentiments represent the more extreme end of the spread of opinion, but they nevertheless highlight the resistance that reformers encountered as they set out to characterise, quantify, and ultimately tackle the challenge of antibiotic misuse. As tensions grew and pressure mounted, the medical community splintered along the familiar fault lines of 'town' and 'gown'. Many doctors in private practice portrayed academic reformers as highfalutin ivory-tower types, divorced from

---

[617] Ibid.
[618] Quoted in Podolsky, *The Antibiotic Era*, p. 123.

the on-the-ground realities faced by time- and resource-strained doctors. 'How many patients should you watch die from lack of prescribing before you give antibiotics?', challenged one physician. 'Those idiots in Washington', he continued, 'should try the practice of medicine for awhile, instead of doing it from the test tube'.[619] In their resistance to reform efforts, physicians borrowed from familiar scripts that pit the atomistic doctor, in touch with the realities of his local community, against out-of-touch forces in the government or the academy encroaching from above. To some, practice guidelines looked less like vehicles for rational medicine than threats to the clinician's authority and autonomy. Inflexible rules delivered from on high could never match the individual judgement of an experienced clinician. The more rhetorically loose among these critics did not hesitate to raise the spectre of 'lysenkoism' or George Orwell's '1984'.[620]

To a limited degree, the MYCIN project experienced similar pushback. In its mildest form, this pushback took the form of ignorance—a failure to appreciate the problem of antibiotic resistance, or this problem's connection to physicians' prescribing practices. The MYCIN team was accustomed to encountering such perspectives in conventional published venues, but occasionally these attitudes surfaced in more surprising (and personal) contexts. In 1977, MYCIN's Victor Yu, on a five-hour, transcontinental airplane ride, found himself seated next to one of the 'site visitors' who had been sent out to Stanford to assess the MYCIN project. Trapped together in a metal tube, Yu and the site visitor began discussing the one thing they had in common: MYCIN. On balance, Yu found the visitor's impressions to be favourable: the visitor reported being impressed by MYCIN's ability to justify itself and by the

[619] Quoted in ibid., p. 124.
[620] Quoted in ibid., p., 127, p. 128-129

team's efforts to systematically evaluate the quality of MYCIN's advice (discussed below). But one line of questions continued to linger in the visitor's mind. As conveyed in the minutes of a MYCIN team meeting, 'His main questions were – "What is the WORTH of an Inf. Dis. Program?" (How often will it be used and for what sorts of consults) and "What is the magnitude of the problem of the incorrect prescribing of antibiotics?"'[621] Clearly, as vigorously as therapeutic reformers campaigned to raise awareness of antibiotic misuse, a sizeable population of otherwise well-informed people remained oblivious to the magnitude of the problem. Encounters like this one prompted Yu to recommend giving all site visitors a bibliography that the MYCIN team had prepared with forty references on the subject of antibiotic misuse.

'HIDDEN IN THE CODE'

Yu's serendipitous meeting at thirty thousand feet, as well as Stanley Cohen's sharp exchange with the National Center for Health Services Research, together illustrate the extent to which ignorance presented a barrier to MYCIN's acceptance. A year later, Yu's colleague William Clancey encountered a different kind of barrier that threatened to be even more difficult to scale: defiance. At an Artificial Intelligence in Medicine (AIM) banquet in 1978, Clancey had a conversation with one physician who seemed to suggest that 'the problem of overprescribing antibiotics, which we argue is an important motivation for the MYCIN program, is not a real problem'.[622] Mystified by the physician's stance, Clancey followed up with him after the banquet. The physician, Harold Goldberger, took this subsequent exchange as an opportunity to

[621] Sherrie Wraith, 'Meeting of March 9, 1977', meeting minutes, printout of email to MYCIN team, 15 Mar. 1977, ESP/10/'MYCIN-Current, 1975-82'.
[622] William Clancey to Harold Goldberger (Shortliffe and Victor Yu copied), email printout, 12 Jul. 1978, ESP/11/'MYCIN – General, 1973-82'.

clarify his position. As it turned out, Goldberger's view was not that antibiotic overprescribing was not a problem; rather, his view was that just as there are risks associated with overprescribing, there are also risks associated with underprescribing. Despite the cries of 'irrationality' coming from antibiotic reformers, there was, Goldberger insisted, a logic—a *rationality*—to physicians' prescribing practices; these clinicians were in the difficult position of having to make medical decisions under conditions of uncertainty, delicately weighing manifold benefits and risks that cut in different directions. As Goldberger explained, 'We can almost never be certain that the therapy we're prescribing is the appropriate one and we have to consider the consequences of erring in various directions. … So we have to consider the relative consequences of prescribing one or several extra antimicrobials that are ineffective against the actual pathogen that's growing and producing its destructive effect in the patient, against the possibility that nothing you're prescribing is going to be effective against the actual infective agent (assuming that an effective drug is available)'.[623]

Goldberger developed these general observations into a more targeted critique of MYCIN. Central to Goldberger's critique was the observation that the developers of MYCIN and practicing physicians were invested in two different, and often conflicting, conceptions of health. Goldberger suggested that Shortliffe and his colleagues, much like leading antibiotic reformers, privileged a population-level conception of health, thus explaining the MYCIN team's concern with the problem of antibiotic resistance. Practicing physicians, by contrast, had no choice but to conceptualise health in terms of the individual—in terms of the person in front of them. Goldberger elaborated on this tension in his response to Clancey: 'Another consideration that I don't think has been

[623] Goldberger to Clancey (Shortliffe and Yu copied), email printout, 17 Jul. 1978, ESP/11/'MYCIN – General, 1973-82'.

fully taken into account is the potential conflict between the interest of an individual patient and the epidemiological concerns that may be *hidden in the code* of a program such as MYCIN. If … one of the concerns of the developers of MYCIN is the proliferation of resistant organisms as a result of overprescription of antibiotics, it would be reasonable for any physician or patient using the program to worry whether the recommendation of the program might be sub-optimal in a particular case because of that concern'.[624] Goldberger's fear that MYCIN's alleged population-level bias might be 'hidden in the code', invisible to the unsuspecting physician-user, is particularly instructive. Such sentiments illustrate how the values of trust, communicability, and transparency informed, often in very fundamental ways, the actions and concerns of both the MYCIN team and its critics.

This discussion of MYCIN—pitting individual-level rationality against population-level rationality—prefigured a line of thinking that would become enduring. Today, deliberations about antibiotic misuse, resistance, and stewardship often centre on the ethical dilemmas surrounding 'the antibiotic commons'. Modern-day antibiotic prescription, many physicians, economists, and ethicists have pointed out, presents a 'tragedy of the commons' scenario, whereby individuals acting perfectly in their own self-interest (or in the self-interest of their individual patients) can endanger the wellbeing of the larger—in this case global—community.[625] Averting the tragedy of the antibiotic commons, many have suggested, demands the cultivation of a new mentality among clinicians, one that foregrounds the health of populations and communities over that of individuals. When it comes to antibiotic prescribing, physicians today are often

[624] Ibid. Emphasis added.
[625] See, e.g., Megan Cully, 'The Politics of Antibiotics', *Nature* 509 (2014): S16-S17.

urged to 'think globally'—to imagine themselves not as a single local actor but as part of a global community confronting a global challenge.[626]

Whatever the legacy of Goldberger's formulation, his point about individual- and population-level health in many ways reproduced the larger town-gown divisions discussed earlier. Indeed, many critics of antibiotic reform efforts came to similar conclusions: while those in the ivory tower, these critics argued, had the luxury of dealing with epidemiological abstractions of health, clinicians in the trenches had the obligation to deal with the wellbeing of the individuals before them. Viewing debates around antibiotic reform through the politics of town and gown goes a long way to making sense of how these reform efforts played out over the latter part of the twentieth century. But such a dyadic framework also elides important nuances. For one, a handful of leading academic clinicians, such as the prominent pharmacologist Louis Lasagna, lent support to so-called town physicians and criticised ideas about regulating prescribing behaviours. What's more, the views of reformers and their critics were far from homogenous. Antibiotic reformers had little trouble agreeing that the present state of antibiotic use demanded some kind of action; agreement about the appropriate shape and scope of that action, however, proved more elusive. Some reformers, such as Finland's ally Calvin Kunin, advocated for clear dictates and national guidelines. Others, such as Finland's Harvard colleague Edward Kass, urged less heavy-handed approaches, favouring educational initiatives that simply augmented physicians' knowledge without curbing their autonomy.[627] From Kunin to Kass, antibiotic reformers arrayed along a spectrum of opinion.

---

[626] Aidan Hollis and Peter Maybarduk, 'Antibiotic Resistance is a Tragedy of the Commons That Necessitates Global Cooperation', *Journal of Law, Medicine & Ethics* 43 (2015): 33-37.
[627] Podolsky, *The Antibiotic Era*, pp. 127, 133-136.

As with the larger national conversations taking place among reformers like Kunin and Kass, the opinions of the MYCIN staff were far from uniform. Different members of the MYCIN project held differing views about the accuracy, universality, and helpfulness of clinical rules and guidelines—and about how rigidly such rules should be applied. These views, however, were not so much expressed outwardly as packaged implicitly in more technical conversations about MYCIN and its inner workings. In the spring of 1976, a debate emerged among MYCIN members around the possibility of rigidly defining what different certainty factors meant. What did it mean, clinically, to say that a certain medical proposition was known with a certainty factor of 0.9 as opposed to one of 0.7? Jan Aikins, one of the MYCIN programmers, outwardly wondered whether defining certainty factors in a standardised manner was, given the inexact state of medical science, impossible:

> [E]ven if we could reach agreement among the infectious disease experts at Stanford as to the "right" CF's to put on our rules, the infectious disease experts on the east coast and other places would probably not agree with us. Now let's take this one step further. Say we are able to assign fairly straightforward meanings to our CF's. Now we have the problem of a doctor in some other part of the country who doesn't want to use MYCIN because our CF's don't agree with what he would use. In other words, by defining our CF's at all rigorously, we're inviting disagreement.[628]

From this discussion of clinical variability and its possible consequences for MYCIN's clinical viability, Aikins ventured that, even if it were possible to define certainty factors with great specificity, such rigid and prescriptive definitions 'might not be such a good idea anyway'. Why not, Aikins implied, give some wiggle room for physicians with varying inclinations and preferences? Aikins' suggestion was consonant with the views of some others in the field of biomedical computing, who argued that

[628] Jan Aikins to 'MYCIN gang', 'On Wednesday's meeting', email printout, 28 Feb. 1976, ESP/10/'MYCIN History, 1972-78'.

computerised decision support tools must 'allow for a spectrum of knowledge and/or opinions instead of just one "golden" standard'.[629] Only by building in such flexibility would it be possible for a system developed at one locale to be adopted at another.

But to other members of the MYCIN team, Aikins' suggestion was preposterous, even heretical. 'I need not remind all of us that we are dealing directly with human lives', Sam Ervin responded, showing little patience. 'If another MD on the east coast disagrees with our CF's and has data (be it strong or weak) as to the basis for his disagreement, then we had better damn well know about it. I claim that one of the advantages of specific criteria for CF's is that it "invites disagreement" (or to put it another way—critical analysis of the rules by non-MYCIN experts is possible)'.[630] According to Ervin, the appropriate response to discordant clinical beliefs is not to throw one's hands up in the air and accommodate both beliefs. No, the responsible course of action is to drill down and figure out which belief is in the right. At stake was nothing less than human lives. This seemingly narrow and technical disagreement between Aikins and Ervin about certainty factors was therefore symptomatic of a much more fundamental disagreement concerning the epistemology of medicine. Whereas Ervin had confidence that there was a 'correct answer' that could be revealed through medical science, Aikins was apparently less sure—or at least less sure that this correct answer could be discerned with much clarity or haste.

Behind these quarrels about certainty factors lay not just philosophical considerations but practical considerations as well. Aikins believed herself to be

---

[629] Jan H. van Bemmel to Shortliffe, email printout, 23 May 1989, ESP/3/'Columbia Presbyterian Med Ctr Retreat, (Arden Homestead, 6/89), 1989'. See also 'Comments on the "Medical Knowledge Base Standard"', ESP/3/'Columbia Presbyterian Med Ctr Retreat, (Arden Homestead, 6/89), 1989'.

[630] Sam Ervin to Aikins and Carli Scott ('MYCIN gang' listserv copied), 29 Feb. 1976, ESP/10/'MYCIN History, 1972-78'.

negotiating a perceived trade-off between precision and viability, accuracy and acceptability. By defining certainty factors rigidly, she argued, the makers of MYCIN would effectively be affixing to their system a ball and chain; they would be limiting how far their expert system could travel. Looser definitions, by contrast, made MYCIN more attractive to a larger and more diverse clinical audience. The notion that computerised systems like MYCIN could be flexible and adaptable to diverse local circumstances was held up as an advantage of new digital technologies over traditional analogue ones. In a 1974 communication to the HEW's Health Care Technology Study Section, written in anticipation of a scheduled site visit to Stanford, Cohen addressed the question MYCIN's 'value added' when compared to the existing information landscape. What advantages did MYCIN have over human infectious disease consultants, or even high-quality infectious disease therapy textbooks?

Cohen enumerated a host of possibilities: a computerised system could respond more nimbly to an ever-evolving base of medical knowledge; it was less prone to embarrass users (computers did not look down upon ignorance in the way human consultants might); it could accommodate uncertain and incomplete medical knowledge (textbooks, according to Cohen, contained only 'well-codified, noncontroversial truths'); it was accessible on demand; it could pool the expertise of multiple infectious disease experts; and it did not tire, applying each rule with the upmost rigour and consistency, whether it was on the first or the one-hundredth consultation of the day. A computerised system also, Cohen argued, promised greater 'local relevance'.[631] 'Textbooks', he noted, 'are designed to transmit general knowledge that is appropriate over a large geographical area, for "typical" patients in various categories…'. MYCIN, by contrast, 'integrates the relevant parts of this general

[631] Cohen to John R. Hall, 14 Mar. 1974, ESP/10/'MYCIN History, 1972-78'.

information for particular patients' and 'applies its knowledge to specific clinical problems which must be dealt with in connection with particular patients'. MYCIN, moreover, could keep track of 'local variations to the general rules within different hospital environments'. With MYCIN, it would thus be possible to accommodate variations in antibiotic sensitivity patterns from one hospital to the next.

Concerns about the extent to which medical knowledge generated at one location might not be generalisable to another location were not unique to the MYCIN team. In fact, other pioneers in computerised diagnosis encountered the challenges associated with transportability and generalisability first-hand: systems based on locally-specific knowledge or data, it was found, failed to travel far outside their place of origin. Francis T. de Dombal's effort to export his computerised system for the diagnosis of acute abdominal pain, AAPHelp, offers a striking example of this phenomenon. When de Dombal and his team at the University of Leeds introduced their Bayesian programme to a hospital in Copenhagen, they watched its accuracy plummet from 80 per cent to 50 per cent, as Marc Berg has described.[632] What went wrong? The team identified a host of issues. Most significantly, the patient population at each site was unique: in one location patients were visiting an emergency department and in the other they were coming to the surgical ward. Whereas any patient could walk into an emergency department, entering a surgical ward required a referral from the patient's general physician. Thus, population statistics derived from patients in an emergency department did not necessarily apply to patients in a surgical ward, and vice versa. Additional cultural, linguistic, and epidemiological differences magnified this problem. Among the symptoms considered by AAPHelp was the extent to which the

[632] Marc Berg, *Rationalizing Medical Work.*

patient's abdominal pain was 'colicky'. But what was the optimal Danish translation for vernacular descriptions of illness like this?

Sometimes the symptoms themselves (not just *descriptions* of symptoms) varied across locations as well. Researchers noted, for example, that the presentation of pancreatitis differed between Copenhagen and Leeds. The difficulties of exporting systems like AAPHelp to diverse clinical environments led de Dombal to conclude that 'databases don't travel'.[633] Optimised for the conditions of one locale, databases (and computerised systems more generally) are often ill-suited to locales that operate under a different set of conditions. As the exchange between Aikins and Ervin suggests, members of the MYCIN team were attuned to the problems of exportability and were committed to avoiding them; they were thinking carefully about the regional specificity of medical knowledge and the possible consequences of this specificity for their enterprise.

Notwithstanding debates about how to calibrate certainty factors, MYCIN was, on balance, highly deferential to the physician's authority and autonomy. There were many possible responses to the problem of antibiotic misuse, from compulsory governmental regulation to optional educational initiatives. Within this spectrum of possible reform policies, MYCIN represented a rather gentler intervention. This deference was not ancillary but fundamental to what MYCIN was and aimed to be. From the outset, concerns about clinical acceptance, autonomy, and authority were never far from the minds of MYCIN's developers. Such concerns even animated many technical decisions about certainty factors and other aspects of MYCIN's inner workings, as the exchange between Aikins and Ervin demonstrates. It would be a mistake to conclude that MYCIN's explanation capabilities were built as a direct

[633] Quoted in ibid., p. 106.

response to wider contemporaneous debates about antibiotic use and clinical autonomy; Shortliffe and his colleagues saw these capabilities as absolutely vital components of medical computer systems, irrespective of their problem domain. But the fact that MYCIN trafficked in the domain of infectious disease—an area of medicine that was, at the time, a site of great tension, division, and resistance—was not unimportant either. This wider context gave the interwoven values of transparency, authority, and trust—values that I argue animated the MYCIN project—even greater weight.

A.I. VS M.D.

Shortliffe and his contemporaries regularly thought about whether the reasoning processes of computerised systems like MYCIN mimicked human forms of reasoning. Beyond such comparisons with respect to styles of reasoning, the Stanford team also compared MYCIN and physicians with respect to the endpoints of reasoning—their final diagnoses and therapeutic recommendations. In their efforts to build trust in MYCIN, Shortliffe and his colleagues pursued something of a two-pronged strategy, although it was not self-consciously seen as such. One prong was creating MYCIN's explanation capabilities—capabilities that aimed to open up the system's recommendations to scrutiny and cede authority to physicians. The second prong was proving, empirically, the accuracy of MYCIN's decisions. Were MYCIN's recommendations more, as, or less accurate than those generated by human physicians? During a period when the clinical trial was ascendant, with new drugs and medical devices held to ever-higher standards of evaluation and validation, the need to assess the system's performance in some empirical, systematic manner was almost a requirement. At first blush, systematic evaluation of MYCIN's decisions seemed like a straightforward enterprise, nothing more than answering the question, 'Was MYCIN

correct, or was it not'. But Shortliffe and his colleagues quickly discovered that behind this seemingly simple question lay deeper complexities. In designing the evaluation study, the MYCIN team had to grapple with complicated questions not just of clinical trial design but also of the philosophy of medicine.

The MYCIN team started to think about how to best go about evaluating the system within a couple years of starting the project. As early as 1974, the team was communicating internally about various procedures for assessing MYCIN's completeness, accuracy, and usefulness. In one such 'Tentative Evaluation Protocol', circulated in March 1974, MYCIN affiliates outlined a two-phased evaluation study.[634] The first phase, to be completed prior to the program's clinical implementation, would tap infectious disease experts at Stanford to 'validat[e] the program's knowledge of bacteremia'. Validation, as outlined in the protocol, would involve running a series of test patient cases on MYCIN. With each patient run, infectious disease experts would examine the same information about the same patient before answering a series of questions about MYCIN's performance:

> (a) What would your therapeutic recommendation have been for this patient?
> (b) What additional questions would you have asked in an effort to better evaluate the patient's status??? Why?
> (c) If your advice differs from that given by the program, do you feel that MYCIN's recommendation is a reasonable alternative regimen?
> (d) Do you feel that the advice given by the program indicates sufficient clinical expertise that the system would be a valuable tool for nonexperts who patients with problems similar to this one???[635]

If the experts concluded that the programme was not running up to par, they would propose modifications to the programme and its rule base. Indeed, MYCIN's

[634] 'Tentative Evaluation Protocol', Mar. 1974, ESP/10/'MYCIN Evaluation, 1974-79'.
[635] Ibid.

developers expected, even hoped, that evaluation would illuminate holes in the system's knowledge base: 'We anticipate that important additional rules will be gleaned by this process', they noted.[636] Evaluation, then, was seen not as a separate step from MYCIN's development; the MYCIN team imagined a constant back-and-forth between developing the system and evaluating it. This sequence—run, evaluate, modify—would be followed iteratively until the programme reached a satisfactory level of performance.

In this first, pre-implementation phase of evaluation, the MYCIN team was interested primarily in MYCIN's accuracy and completeness. Before physicians could be expected to adopt MYCIN, they needed reasonable proof of the integrity of the system's recommendations and the knowledge upon which such recommendations were based. In the second, post-implementation phase of evaluation, a different set of questions would be operative. To demonstrate the appropriateness of MYCIN's recommendations was not to prove that the system, once implemented, actually improved either physician behaviours or patient outcomes. As the tentative evaluation protocol explained, 'Since a major project goal is to improve care of patients with infectious diseases, not simply to persuade physicians to follow the advice of a computer program, we must also evaluate the impact of the program upon patient well-being'.[637] Surrogate endpoints would not do. The team, originally, wanted a clear understanding of how the system affected patient morbidity and mortality. The kind of evaluation that the protocol's authors had in mind involved appraising (1) outcomes of patients treated for bacteraemia and (2) bacteria antimicrobial prescribing habits of physicians two to three months *prior* to the system's introduction in the Stanford

[636] Ibid.
[637] Ibid.

Hospital and then appraising these same metrics again after several months of on-going MYCIN use.

As the MYCIN team thought through the logistics of evaluation, they identified a number of key sticking points. Most importantly, they were attempting to evaluate something for which there was no single, self-evidently 'correct' answer. In medicine, the team realised, often multiple therapeutic approaches can be justified; the reason for choosing one approach over another can be as much a matter of values as of empirics. Medical practice accommodates a host of *different* values and rationalities, and the 'superior' therapeutic approach is often a consequence of how one opts to weigh these co-existing values, interests, and considerations. No clear-cut laboratory test can be ordered to tell one how to best assign these weights. Shortliffe was acutely aware of these challenges, and he went out of his way to correct colleagues who failed to appreciate the complexities of clinical medicine. In reviewing a manuscript submitted to the *Annals of Internal Medicine*, Shortliffe critiqued the paper for its assumption that all medical knowledge could be sorted neatly into one of two buckets: 'The discussion seems to assume that knowledge is either "right" or "wrong," yet people in the field are aware that absolute gold standards of performance are difficult to define'.[638] Buchanan similarly acknowledged how it could be 'difficult to define precisely the term "appropriate therapy"'.[639] Appropriate by what criteria, and by whose standards?

Even outside the realm of values, the MYCIN team discovered that medical decisions were difficult to evaluate in a binary, right-wrong fashion for a more practical reason: different experts believed different things. The MYCIN team commented on

---

[638] Shortliffe, manuscript review for *Annals of Internal* Medicine, 9 May 1984, ESP/2/'Pittsburgh, Univ. of, 1980-95'.
[639] Gregory Freiherr, 'MYCIN: A Medical and Educational Aid', *Research Resources Reporter* (US Department of Health, Educaiton and Welfare, Public Health Service, National institutions of Health), December 1979, Volume III, no. 12., pp. 1-6, p. 2.

this lack of consensus soon after starting work on the project. In a 1972 memo, Stanley Cohen noted, 'In certain instances there might be so much difference of opinion that no reasonable hierarchy of recommended therapy can be defined'.[640] By 1977, Cohen and his colleagues were no longer talking about expert variability in a speculative manner—in the realm of 'mights' and 'maybes'. After five years of work, preliminary findings had confirmed Cohen's earlier suppositions. As he reported in a 1977 letter, 'Our group at Stanford has been studying clinical decision making in infectious disease therapy for the last several years. As you may surmise, we have observed considerable variability among different infectious disease specialists in their antimicrobial recommendations when presented with identical patients. Part of this variability', they reasoned, 'stems from level of training and from geographic and institutional biases'.[641] The challenges presented by expert variability were twofold. One, it smothered hopes of evaluating MYCIN's performance against a single, gold standard. And two, it meant that even if the team designed a study that evaluated MYCIN's performance against that of certified experts, the team would have to somehow account for disagreements among the experts themselves.

These were significant challenges, to be sure, but the MYCIN team viewed them as surmountable. '[T]he real world problem area', Shortliffe explained in a 1975 letter, 'is filled with disagreement and it was thus a challenge to try to develop an evaluation mechanism that would control for such lack of unanimity'.[642] But try the MYCIN team did. The team published their first evaluation study—of the system's

---

[640] Cohen to Thomas Merigan and Stanton Axline, memo, 'Proposed computer project involving antibiotics therapy', 25 Feb. 1972, ESP/10/'MYCIN History, 1972-78'.
[641] Yu and Cohen to Morton N. Swartz, 9 Aug. 1977, ESP/4/'Victor's Letters, 1977-78'.
[642] Shortliffe to William Mark, email printout, 21 Jul. 1975, ESP/10/'MYCIN, 1973-77'.

bacteremia knowledge base—in a 1979 issue of *Computer Programs in Biomedicine*.[643] In the study, the MYCIN team sent on fifteen patients to ten experts, five infectious disease specialists from Stanford and five infectious disease specialists from other institutions. The experts then judged whether MYCIN's decision-making process and recommendations were identical or acceptable alternatives to their own recommendations. By this method of evaluation, MYCIN met Stanford's standards of acceptable practice in over 90 percent of cases. This bacteraemia study, however, left an important question yawning open before the MYCIN team. Where MYCIN and the physicians differed in their recommendations, who was in the wrong: the physicians, MYCIN, or both?

The team aimed to address this critical shortcoming in a second evaluation, this time of MYCIN's meningitis, rather than its bacteraemia, knowledge base. On 9 August 1977, Morton Schwartz, an infectious disease expert at MGH, received an unsolicited letter from MYCIN's Victor Yu and Stanley Cohen: 'In order to further study this phenomenon [of clinical variability in antibiotic prescribing], we are asking for the participation of a group of distinguished academic infectious disease experts who have done clinical studies involving patients with meningitis. We would like you to look over 10 clinical summaries of patients with meningitis and complete a short questionnaire concerning antimicrobial recommendations'.[644] Schwartz was just one of at least eight such experts to receive the invitation. What Schwartz and the other

---

[643] Victor L. Yu, Bruce G. Buchanan, Edward H. Shortliffe, Sharon M. Wraith, Randall Davis, A. Carlisle Scott, and Stanley N. Cohen, 'Evaluating the Performance of a Computer-Based Consultant', *Computer Programs in Biomedicine* 9 (1979): 95-102.
[644] Yu and Cohen to Swartz, 9 Aug. 1977. The experts the MYCIN considered reaching out to included: Morton Swartz (Harvard), George McCracken (Parkland), Jay Sanford (United Health Services), Philip Lerner (Case Western), Stuart Levin (Chicago), Ralph Feigin (St. Louis), Robert Petersdorf (University of Washington), John E. Bennett (NIH), Philip Dodge, Donald Armstrong (Cornell), James Rahal (NYU).

experts were not privy to was the real purpose of the study: they were not informed that one of these ten 'clinical summaries' was of *particular* interest to the researchers. This summary had been produced by computer—by MYCIN. To serve as points of comparison, the other nine clinical summaries had been generated by medical personnel with varying levels of expertise and training.

In this letter to Schwartz, one begins to see the outlines of the ultimate methodology behind what would become the most rigorous evaluation of MYCIN. Victor Yu proposed the evaluation methodology in a 1976 email to his colleagues:

> Pick a selected number of interesting and/or challenging cases, 8 to 12 in number. … Have a 3rd year medical student, 4th year medical student, intern, resident, ID fellow, ID staff evaluate each case selecting his treatment recommendations, if any. Have 2 or 3 ID staff who are nationally recognized experts in Meningitis give their treatment recommendations. Then, they would "grade" the treatments given by our volunteers (students, housestaff, fellows, Stanford staff) plus MYCIN without knowing the identity of each volunteer (ie human vs MYCIN). Agree, Disagree, and Close call would be used.[645]

The entire purpose of this methodology—pooling MYCIN's recommendations together with the recommendations of various medical personnel and then subjecting that pool of recommendations to a second layer of evaluation—was 'to achieve just the comparative power that the first two [evaluation studies] lacked'.[646]

The evaluation as ultimately published hewed closely to Yu's proposal.[647] It had two phases. In the first phase, a physician, with no affiliation to the MYCIN project and with little-to-no knowledge of the system's capabilities, selected ten patients with

---

[645] Yu to Axline, Shortliffe, Cohen, Wraith (Scott, Fagan, van Melle, and Clancey copied), email printout, 29 Jan. 1976, ESP/5/'Bacteremia MYCIN – Patient Runs, 1975'.
[646] <LETSINGER>EVAL.DRAFT;2, unpublished printout, 28 Nov. 1979, ESP/10/'MYCIN Evaluation, 1974-79'.
[647] For a more detailed draft protocol for the evaluation study, see 'The meningitis evaluation', unpublished printout, 4 Apr. 1977, ESP/5/'Bacteremia MYCIN – Patient Runs, 1975'.

infectious meningitis according to predetermined criteria.[648] These patient cases were compiled into detailed summaries, which contained information concerning the patients' medical history, physical examination, laboratory data, and the immediately prior hospital course. These detailed summaries were then sent to different medical personnel, representing a spectrum of expertise, for evaluation: a medical student, a medical resident, a senior postdoctoral fellow in infectious disease, five faculty members in Stanford's Division of Infectious Disease, and MYCIN. Using only the information contained in the patient summaries, each prescriber reported what he, she, or it determined to be the optimal antimicrobial therapy regimen for each of the ten cases.

Having generated a pool of recommendations that could be compared and judged, the evaluation entered its second phase. The MYCIN team sent the ten patient summaries together with the resultant therapeutic recommendations to eight 'super experts'—people, like Schwartz, who were deemed experts in the management of meningitis and had published case reports dealing with the subject. The super experts were asked to review the patient cases and come up with their own prescriptions, before evaluating the therapeutic recommendations of the original prescribers. The super experts rated the original prescribers in a tripartite fashion: the prescription could be (1) 'equivalent' to that of the super expert, (2) different from that of the super expert but nevertheless 'acceptable', or (3) 'not acceptable'.

[648] This physician was Robert Blum. On some early discussions of possible selection criteria, as well as a more thorough discussion of the goals of the study, see 'Goals for the meningitis evaluation', unpublished printout, 12 Apr. 1977, ESP/5/'Bacteremia MYCIN – Patient Runs, 1975'.

The results of the study, published in *JAMA* in 1979, looked pretty good for MYCIN.[649] The system's recommendation was deemed either 'equivalent' or 'acceptable' in sixty-five percent of the evaluations. While such a result surely left room for improvement, it compared favourably to the other prescribers evaluated in the study. Medical students issued acceptable prescriptions only 30 percent of the time—an outcome that residents bested by 15 percent. Even the five members of the Stanford faculty, some of whom were published experts in the field of infectious disease, sunk below MYCIN's level of performance, offering an appropriate recommendation in between 42.5 and 62.5 percent of cases.

In the published study, the authors offered four reasons to make sense of what they characterised as MYCIN's 'strong performance', especially when compared to the performance of human clinicians.[650] First, MYCIN simply 'knew' more; the system's knowledge base had been constructed over the course of many years, drawing on not just the knowledge of dozens of infectious disease experts but also the information contained in the medical literature and the medical records at three hospitals. Second, the team argued that MYCIN was immune to the lapses of memory that plagued so many human physicians: 'A popular maxim among physicians is that "one has to think of the disease to recognize it." This is not a problem for the program; rare diseases are never "forgotten" once information about them has been added to the knowledge base, and risk factors for specific meningitides are systematically analyzed'.[651] Third, the system applies understandings of diagnostic importance consistently and

---

[649] Victor L. Yu, Lawrence M. Fagan, Sharon M. Wraith, William J. Clancey, A. Carlisle Scott, John Hannigan, Robert L. Blum, Bruce G. Buchanan, and Stanley N. Cohen, 'Antimicrobial Selection by a Computer: A Blinded Evaluation by Infectious Disease Experts', *JAMA* 242, no. 12 (1979): 1279-1282.
[650] Ibid, p. 1282.
[651] Ibid.

systematically; it does not 'jump to conclusions on the basis of an isolated finding nor does it neglect to ask for key pieces of information'.[652] Finally, the system's knowledge is current, regularly modified in light of new findings.

For those already inclined towards computerised medicine, the study heralded a bright future. One Toronto-based laboratory technologist cited the MYCIN evaluation study (among others) as evidence that a digital revolution in medicine lurked just over the horizon. Immediately after summarising the result of the MYCIN evaluation study, he concluded, 'Computer applications in medicine will come of age in the 1980s. Despite resistance of many doctors to the use of these machines, the computer will lead to a new era in medical care. With the cooperation of physicians, computers may become to the medical system what the wheel was to mankind'.[653] Yet, when the 1980s came, MYCIN remained confined to experimental settings, and Shortliffe and his Stanford colleagues had started to shift their attention to a different medical programme. To show that a computerised diagnostic system could 'beat' its human counterparts in controlled experimental settings was a far cry from demonstrating the system's use in the real world.

A 'GLORIFIED ALGORITHM'?

After a decade's worth of effort, the firewall between the *research* side of medical AI and the *clinical* side of medical AI remained as strong as ever. Shortliffe and his colleagues were eager to break through. Shortliffe offered a sobering assessment of the situation in a 1982 memo to his colleagues at the Heuristic Programming Project (HPP): 'Despite active research into the development of medical

[652] Ibid.
[653] Richard S. Jones, 'Letters: Use of Computers in Medical Practice', *Canadian Medical Association Journal* 122, no. 4 (1980): 400.

consultation systems over the last decade, no system from the AIM community has yet been introduced for ongoing use in clinical setting. Although many of the most important theoretical questions involve the optimal design of the reasoning component of such system, the development of "useable" systems that will be accepted by physicians also requires attention by the AIM research community'.[654]

In MYCIN's case, this distance from a real clinical setting was especially tragic. Many of the system's most prized features concerned its usability and acceptability within real world clinical environments. Because MYCIN remained a research rather than a clinical enterprise, these features were never meaningfully tested. In branching out from the MYCIN project, Shortliffe and his colleagues thus prioritised the selection of a new domain that would allow them to fast-track a computerised consultation system into the actual clinical setting.

The team quickly settled on oncology protocol management and developed a system called ONCOCIN. In the 1980s, as is the case today, many treatments in oncology were experimental. For many types of cancers, physicians and researchers did not have a clear sense of the comparative benefits and toxicities of different treatment regimens and modalities. As a result, oncologic treatment often involved enrolling patients into different experimental treatments, with the short-term hope of curing patients and the long-term hope of enriching medicine's understanding of comparative efficacy. Elaborate protocols—often spanning dozens or even hundreds of pages in length—were created to define the particulars of a given treatment plan, set certain eligibility requirements, and much else. Strict adherence to the protocol was a must: because the protocols were not only treatment descriptions for individual patients but

[654] Shortliffe to HPP Faculty and Staff, 'ONCOCIN: An Oncology Protocol Management Consultation System', memo, 16 Feb. 1982, ESP/2/'ONCOCIN 1977-89 (fl 1 of 2)'.

also key components of a larger experimental design, departure from the protocol could endanger the interpretability and analysis of treatment results.

The team also hoped that in selecting the domain of oncology protocol management they would be filling an unmet clinical need. Oncology protocols could be extraordinarily complex, delineating not only the appropriate course of treatment but also the appropriate management of countless events, problems, or complications that could arise during the treatment. As the group wrote in 1980, 'Because protocols may sometimes be very complex, it can be difficult for busy physicians to remember the details sufficiently well to comply with the protocol adequately'.[655] The limitations of the human mind, combined with the hectic working environment of an oncologic clinic, nearly guaranteed deviation from protocol. The team further explained in 1984: 'Although an effort is made to have the documents [protocols] available in the oncology clinics when patients are being treated for their tumors, it is often the case that a busy clinic schedule, coupled with a complex protocol description, leads a physician to rely on memory when deciding on drug doses and laboratory tests. Furthermore, solutions for all possible treatment problems cannot be spelled out in protocols. Physicians use their own judgment in treating these patients, resulting in some variability in treatment from patient to patient'.[656] ONCOCIN might thus help cognitively- and administratively-overstretched physicians adhere more closely to protocols of ever-increasing complexity.

---

[655] <CLANTON>ONCOCIN.INTERFACE;1, unpublished report (computer printout), 6 Feb. 1980, ESP/2/'ONCOCIN – Logistics, 1978-83'.

[656] Edward H. Shortliffe, A. Carlisle Scott, Miriam B. Bischoff, A. Bruce Campbell, William van Melle, and Charlotte D. Jacobs, 'An Expert System for Oncology Protocol Management', in Bruce G. Buchanan and Edward H. Shortliffe (eds.) *Rule Based Expert Systems* (Reading, MA: Addison-Wesley, 1984), 653-668, pp. 654-655.

But even beyond perceptions of clinical need, oncology protocol management gained favour among the team for a more self-interested reason: it was deemed supremely computerisable. 'ONCOCIN's domain was selected', Shortliffe clarified in his 1982 memo, 'precisely because its knowledge base is unusually well-formalized and highly structured. We wanted to develop a system that could be implemented for ongoing use in the short term, thereby allowing investigation into human issues that cannot be readily address when a system exists solely in a research environment'.[657] What would physicians do when confronted with a computerised consultation system? How might such a system disrupt, or improve, a clinic's workflow? The domain of antibiotic selection offered a nice playground for the invention of a new kind of computerised system—the rule-based expert system. It would be the domain of oncology that would allow for the introduction of such a system in a real clinical setting. Their domain selected, the team began encoding a series of protocols for Hodgkin's and non-Hodgkin's lymphoma in the ONCOCIN system. Their selection of lymphoma protocols was largely circumstantial: these were protocols that had been developed at Stanford in its capacity as a major lymphoma treatment centre nationally. ONCOCIN was born. In May 1981, ONCOCIN went live at Stanford's Debbie Probst Oncology Day Care Center.[658] Three mornings per week, eight fellows in oncology used the system—implemented on a DECSYSTEM-20 computer—to help manage patients enrolled in lymphoma chemotherapy protocols.

ONCOCIN's debts to MYCIN were clear. It was, after all, a rule-based system, containing hundreds of IF-THEN rules that would be immediately recognisable to

[657] Shortliffe to HPP Faculty and Staff, 'ONCOCIN: An Oncology Protocol Management Consultation System', 16 Feb. 1982.
[658] ONCOCIN Annual Report, unpublished report, n.d. (circa late 1982), p. 4, ESP/2/'ONCOCIN 1977-89 (Fl 1 of 2)'.

anyone with a passing familiarity with MYCIN. Yet the computerised system that the team eventually developed differed from MYCIN in important ways. Not long after building MYCIN, its developers realised that the system's reasoning framework might be useful not just for infectious disease but for other domains as well. This reasoning framework was created by extracting the infectious disease knowledge base from the expert system, leaving behind the empty reasoning shell. The team would come to call this reasoning shell E-MYCIN, for 'Empty MYCIN'. By replacing the infectious disease knowledge base with IF-THEN rules about a new domain, other researchers could quickly create new MYCIN-like expert systems to assist in myriad new problems. Before long, (E-)MYCIN spawned a number of spin-offs: SACON was created to help engineers analyse how buildings and other physical structures would respond to varying loads; CLOT, to assist physicians determine the presence of blood clots; and PROSPECTOR, to direct where, geographically, one could profitably drill for mineral deposits. (In 1980, PROSPECTOR proved itself commercially viable, with the accurate prediction of a large molybdenum deposit in Washington state.)[659]

ONCOCIN, however, was different from these other MYCIN progeny: creating ONCOCIN involved more than plugging the knowledge base of oncology protocol management into the E-MYCIN shell. The team reprogrammed ONCOCIN 'from scratch'.[660] This reinvention, the team felt, was necessitated by important differences between the domains of infectious disease and oncology protocol management. For example, whereas MYCIN only considered a patient's situation at a single moment in

[659] Robert S. Engelmore, 'Artificial Intelligence and Knowledge Based Systems: Origins, Methods and Opportunities for NDE', in D.O. Thompson and D.E. Chimenti (eds.), *Review of Progress in Quantitative Nondestructive Evaluation*, Vol. 6 (Boston: Springer, 1987), 1-20, p. 13.
[660] In fact, 'portions of the E-MCYIN code' were recycled in the creation of ONCOCIN. 'Longer Progress Report, Project 1 (1979-1980)', unpublished report, n.d., ESP/2/'ONCOCIN 1977-89 (Fl 1 of 2)'.

time, ONCOCIN required serial consideration of a patient's evolving situation over many months.

The commonalities between ONCOCIN and MYCIN still outweighed their differences. Above all, the two shared an emphasis on explanation. The ONCOCIN system had two primary subsystems: the Reasoner, 'a rule-based expert consultant that is the core of the system', and the Interviewer, 'an interface program that controls a high-speed terminal and the interaction with the physicians using the system'.[661] The group felt strongly that ONCICIN's success hinged equally on the Reasoner and the Interviewer: a smart computer system that could not compellingly engage with clinicians may as well be dumb. As related in an early progress report on the project, 'The actual mechanics of computer terminal interaction is as important to a clinical system's acceptance as the quality of the program's advice'.[662] The team focused on three aspects of the physician-computer interface. First, the team worked to ensure that ONCOCIN's hardware was optimised for speed, with the selection and instalment of high-speed CRT terminals. Second, they aimed to reduce the burden caused by excessive typing requirements. Initially the team wanted to create a light-pen or touch-screen interface, but a combination of financial and technological limitations steered them instead to develop a customised 21-character keypad. Third, and most importantly, they invested in the system's explanation capabilities: '[W]e wanted to maintain the explanation and justification capabilities which we have argued are crucial to the acceptance of clinical consultation systems'.[663]

---

[661] Shortliffe et al., 'An Expert System for Oncology Protocol Management', p. 653.
[662] 'Summary, Project 1: ONCOCIN', unpublished report, n.d., ESP/2/'ONCOCIN 1977-89 (Fl 1 of 2)'.
[663] 'Longer Progress Report, Project 1 (1979-1980)'.

As with MYCIN, ONCOCIN's user interface, and its explanation capability in particular, encoded certain ideas about the physician's identity and authority. '[T]he physician generally has a rather self-sufficient, active, controlling personality', explained a 1980 report on ONCOCIN's user interface. 'Thus, an unwanted intrusion on his lifestyle or environment is likely to be strongly resisted. Any machine that is designed to help him must be clearly and readily under his control'.[664] Sometimes these ideas about the physician's identity and authority were couched in more gendered terms. As the same report noted, 'While the issue of security and confidentiality is critical, so is the issue of friendliness to the physician. A system that is hard to get to use, that does not quickly and readily recognize the physician as does the secretary and nurses, is seen as hostile, harsh, and unpleasant'. ONCOCIN's 'interface program' thus contained cultural ideas about the identity of the physician—his preferences, tendencies, and even personality ('self-sufficient, active, controlling').

The purpose of ONCOCIN's 'Interviewer' was not simply to explain the system's reasoning. It was also to perform deference. One of the ONCOCIN project members, Cliff Wulfman, described how '[t]here are two ways to view Oncocin's relationship to the physician using it. One way is to envision Oncocin as a Domain Expert Consultant. In this view, the Reasoner predominates; IT asks the questions, IT produces recommendations, and everything else is side effect. The other view (and the preferred one, I believe) makes Oncocin a Helpful Assistant, looking over the physician's shoulder as he fills out forms and decides what therapy to give on this visit. … In this view, the Interviewer is not an "interviewer" at all; it is an archiving tool'.[665] Here again, important to note is how the 'Interviewer' components of ONCOCIN

[664] <CLANTON>ONCOCIN.INTERFACE;1.
[665] Cliff Wulfman, 'The Interviewer as Input Device and Archivist's Tool', n.d. (circa 1985), ESP/7/'ONCOCIN, 1979-88'.

shaped the ontology of the system: ONCOCIN's ability to explain itself and to engage deferentially with the physician transformed its status from that of 'Domain Expert Consultant' to that of 'Helpful Assistant'. Whether the system was conceived of as the former or the latter, its designers believed, would greatly influence its clinical acceptance.

They were right. Their claim that computerised systems were purveyors of 'advice' rather than 'dogma' resonated with concerned physicians. In 1987, AMA representative Daniel K. Harris wrote an editorial for *JAMA*, describing both the opportunities and the challenges introduced by new strides in the computerised decision support field.[666] In a letter to the editor the University of Chicago physician James Levin characterised Harris' piece as expressive of 'the fear that physicians may have that, if not vigilant, they will either abdicate, or somehow be forced to relinquish control of their diagnostic and therapeutic decision-making to computer software'.[667] Yet Harris was reassured by Shortliffe's claim that computerised systems would remain under the dominion of the physician. Harris wrote, 'Shortliffe expresses the same consideration: "…Decision-support programs are intended to serve as *tools* for the trained practitioners, who retain ultimate responsibility for determining diagnostic and therapeutic strategies."'[668] Having marshalled this Shortliffe quote, Harris had little else

---

[666] Daniel K. Harris, 'Computer-Assisted Decision Support: A New Available Reality', *JAMA* 258, no. 1 (1987): 86.
[667] James E. Levin, 'Computer Programs to Support Clinical Decision Making', *JAMA* 258, no. 17 (1987): 2375. For other articles in conversation with these articles, see Edward H. Shortliffe, 'Computer Programs to Support Clinical Decision Making', *JAMA* 258, no. 1 (1987): 61-66; Edward H. Shortliffe, 'Computer Programs to Support Clinical Decision Making – Reply', *JAMA* 258, no. 17 (1987): 2375-2376; G. Octo Barnett, James J. Cimino, Jon A. Hupp, Edward P. Hoffer, 'Computer Programs to Support Clinical Decision Making – Reply', *JAMA* 258, no. 17 (1987): 2376; Daniel K. Harris, 'Computer Programs to Support Clinical Decision Making – Reply', *JAMA* 258, no. 17 (1987): 2376.
[668] Harris, 'Computer Associated Decision Support', p. 86.

to add beyond an 'Amen': 'As we take advantage of these exciting technologies, it would be well to heed these words'.[669]

Even before ONCOCIN's rollout, Paul Chang, a member of the ONCOCIN team, conducted a series of pre-implementation interviews with physicians to better understand their views on computerised systems like ONCOCIN. Their responses covered a wide range. Some interviewees expressed excitement about the possibility of a computer lightening their administrative burden. Others expressed doubt or concern. As related by Chang, 'A few interviewees seriously doubted the ability for "computer algorithms" to effectively deal with the "art of medicine" – that it is impossible for a computer program to deal with situations that cannot be answered simply by "flowcharts or simple rules." Many questioned whether or not a program can ever get an "adequate picture ('gestalt') of a patient's state of health by the mere asking of a few questions"'.[670] The interviewed physicians also touched upon issues of understandability and trust. At least one physician emphasised that the system's explanations must be 'intuitively reasonable' to the user while also avoiding 'Bayesian/probabilistic magic'.[671] For some, the importance of transparency extended beyond the mere medical knowledge within ONCOCIN and to the *kind of people* who created the system in the first place. Were these people credible? Did they understand the on-the-ground realities of clinical practice? As Chang characterised one discussion, 'One must also emphasize "human input" into the system: the user must be convinced that "sympathetic and experienced" physicians were involved in the design and implementation of the system'.[672] The system, one physician noted, needed to be more

[669] Ibid., p. 86.
[670] Paul Chang, 'Pertinent Findings from Physician Test Interviews', unpublished report, n.d., ESP/2/'[ONCOCIN], 1979-83'.
[671] Ibid.
[672] Ibid.

than a 'glorified algorithm'. Concerns about trust and transparency shaded into concerns about control and authority. Chang summarised the view among some of his subjects that '[p]hysicians much be able to always override the advice given by the computer—"the physician must always be in charge."'[673]

For all that the ONCOCIN team had learned from their earlier experiences with MYCIN—and from their pre-implementation interviews with physicians and oncology fellows—they nevertheless encountered certain unanticipated (albeit interesting) challenges. One problem that the ONCOCIN team encountered derived from a phenomenon they had come to know well during their work on MYCIN: clinical variability. An implicit assumption made in the first iteration of ONCOCIN was that heterogeneous medical decisions—those that deviate from the protocol—were significant either for patient care or for research purposes. Indeed, clinical variability was the enemy of the protocol. As one team member wrote in 1980, 'The purpose of a protocol is uniformity of care so that different treatment regimens can be compared more readily'.[674]

However, the ONCOCIN team came to realise that, in practice, physician and protocol can differ in ways that were more stylistic than substantive. In these cases, physicians bristled at having to override a computerised system—and at having to provide a justification for doing so. The team elaborated on this realisation in one of their annual reports:

> The second problem was unanticipated until the system was implemented, but it is having a major impact on our current work. ONCOCIN provides a therapy recommendation for the physician but allows those recommendations to be overridden if the physician prefers an alternative treatment. If the fellows change the therapy, they are asked to indicate the reason for preferring their therapy to what ONCOCIN suggested. When there is a major disagreement between the fellow and the computer, the physicians seldom object to indicating

[673] Ibid.
[674] <CLANTON>ONCOCIN.INTERFACE;1.

> the reasons for the change, and their entries have been useful in further debugging the program's knowledge of lymphoma treatment. However, the changes they make are frequently minor ones that are more stylistic than substantive (e.g., some physicians prefer always to give prednisone in multiples of 20 mgs when they are dealing with large doses; others prefer to be more precise in their dosing and will ask patients to take combinations of 5 and 20 mg tablets). In these cases, the physicians tend to be annoyed by requisitions for reasons that they have disagreement with ONCOCIN; they do not see the slight adjustments as being real disagreements.[675]

The discovery of 'stylistic' differences among physicians highlighted certain tensions inherent in ONCOCIN's task. Aggressive eradication of clinical idiosyncrasy by computerised protocol might reduce clinically significant heterogeneity. But it also risked butting up against the physician's very identity—the personalised styles of practice with which many physicians came to identify.

Such realisations during the process of clinical implementation drove revisions to the system. Some of these revisions were geared toward identifying—and respecting—the 'stylistic' forms of decision making and treatment. The ONCOCIN team began 'developing criteria for determining when a physician's preferred treatment is a clinically significant difference from what ONCOCIN has recommended'.[676] Recognising the conceptual distinction between 'substantive' and 'stylistic' modes of medical practice was one thing; delineating the criteria that would actually differentiate between the two was another entirely. Though such a task was far from straightforward, the ONCOCIN team recognised its importance: 'The knowledge for determining what changes are "important" is different from the kind of knowledge in the protocols themselves. It is more "judgmental" and imprecise, but seems to be an important addition to the program's knowledge base'.[677]

---

[675] ONCOCIN Annual Report, p. 6.
[676] Ibid.
[677] Ibid.

Treatment decisions in medicine rest not just on medical evidence but also on the patient's values, goals, and circumstances. Yet it was impossible for the ONCOCIN team to formally encode all the possible social, moral, and personal drivers of therapeutic decision making. This hole in ONCOCIN's knowledge base led the system to deliver recommendations that—while justifiable on the basis of oncologic protocol and medical evidence—were 'wrong' considering the wider context of patients' life circumstances. Shortliffe recounted one such case in 1987. On a typical Friday afternoon, a young woman with lymphoma was seen at the Stanford oncology clinic. ONCOCIN recommended business as usual: full, 100 percent continuation of the patient's current chemotherapy regimen. The physician, however, sharply disagreed, opting instead to give her no treatment. When the clinic data manager encountered this discrepancy, she pressed the physician: what, she probed, could account for such a stark disagreement? As Shortliffe later described it, 'the physician explained that the patient was to be a bridesmaid the next day and had asked to delay her treatment (which routinely caused a day or so of nausea) until after the weekend'.[678] Shortliffe recognised that such a decision required knowledge that would be difficult, if not impossible, to incorporate into a computerised system like ONCOCIN: 'Her request was eminently reasonable, but formulating a proper response required both medical expertise (to determine if a brief delay in treatment would be dangerous) and world knowledge (about bridesmaids, the nature of nausea, and why a patient would not want to be nauseated during a wedding)'.

Other post-implementation changes to the programme continued along similar lines. Many proposed revisions aimed to shift ONCOCIN's role from one of active

[678] Edward H. Shortliffe, 'Computer Programs to Support Clinical Decision Making', *JAMA* 258, no. 1 (1987): 61-66, p. 66.

guidance during the decision making process to one of passive verification of already-made decisions. Under these changes, ONCOCIN would be in the business of *assessment* rather than *management*: '[T]he observation that physicians frequently believe they know what to do and would rather not be bothered by the program's suggestions, have led us to develop a prototype modified version of ONCOCIN that verified the physician's plan rather than offered a therapy recommendation for every case. We have referred to this new work as "hypothesis assessment"'.[679] The ONCOCIN team went great lengths to qualm physicians who disliked being told what to do by a computer. The physician 'owned' the medical decision; it was critical that ONCOCIN's participation in the decision making process reflect this. '[W]e have continued', the team explained, 'to develop modifications to ONCOCIN that will permit it to function as an "observer" of the physician's own decisions rather than as a primary source of advice'.[680]

ONCOCIN's creators framed this as a shift towards a 'critiquing model' of decision support. One member of the team elaborated on this reframing during a 1984 demo of ONCOCIN's capabilities:

> Now I'd like to describe how ONCOCIN has been adapted to critique physician's therapy plans. We made this adaptation because feedback from physicians who use the system indicated that a system that can accept and critique proposed therapy plans may be more acceptable to them. Because ONCOCIN is used for all visits that lymphoma patients make to the oncology clinic, physicians often use the system for routine cases about which they might not normally seek advice. By allowing the physician to suggest treatment first, we eliminate the need for the physician to "override" ONCOCIN's recommendation when it contains only minor differences. Instead, the critiquing model detects any differences between the physician's proposed plan and its own. If those differences are minor, *no interruptions need occur*, and ONCOCIN can accept the physician's recommendation. If significant differences exist, those differences are pointed out, and a critique is offered which consists of explanations [for] why ONCOCIN concluded what it did. The critique process consists of two significant computation steps after the physician

[679] ONCOCIN Annual Report, p. 6.
[680] Ibid., p. 7.

> has entered the patient data and the proposed recommendation. First, the proposed recommendation must be *evaluated*. Then any clinically significant differences between the two are *explained*.[681]

For a period of time, ONCOCIN generated promising results. Shortliffe and his colleagues found that ONCOCIN promoted increased completeness in clinical data collection, with physicians collecting and reporting the data expected by oncology protocols at higher rates following ONCOCIN's implementation.[682] Beyond promoting thoroughness of recordkeeping, the advice that ONCOCIN dispensed, the team discovered, also met reasonably high standards. In a study of 415 clinic visits involving ONCOCIN, the team found that the computer system's recommendations were identical to the treatment provided by physicians in 46 percent of cases. Where physicians and the system disagreed, an independent panel of expert oncologists found the differing recommendations to be equally acceptable.[683] Despite its promise, ONCOCIN never caught on widely.

* * *

The MYCIN project's legacy did not end with ONCOCIN. Even as its medical instantiations stumbled, its other descendants marched onward, reaching fields far outside that of medicine. The large demand of MYCIN's software inspired Edward Feigenbaum to found a company, Teknowledge, whose goal was 'to migrate EMYCIN

---

[681] ONCOCIN demo talk, unpublished notes, Feb. 1984, ESP/2/'ONCOCIN 1977-89 (Fl 1 of 2)'. The author of these notes is not listed. It is probably Shortliffe.
[682] Daniel L. Kent, Edward H. Shortliffe, Robert W. Carlson, Miriam B. Bischoff, Charlotte D. Jacobs, 'Improvements in Data Collection Through Physician Use of a Computer-Based Chemotherapy Treatment Consultant', *Journal of Clinical Oncology* 3 (1985): 1409-1417.
[683] David H. Hickam, Edward H. Shortliffe, Miriam B. Bischoff, A. Carlisle Scott, and Charlotte D. Jacobs, 'The Treatment Advice of a Computer-Based Cancer Chemotherapy Protocol Advisor', *Annals of Internal Medicine* 103, no. 6 (1985): 928-936.

into the commercial domain, make it industrial strength, sell it, and apply it'.[684] MYCIN's influence even extended beyond the successes of its immediate spin-offs. The project defined a new paradigm of research and practice—the rule-based expert system—that would dominate the field of AI through the 1980s. The AI pioneer Allen Newell characterised the magnitude of this shift in 1984: 'The last seven years have seen the field of artificial intelligence (AI) transformed. This transformation has been generated by the emergence of *expert systems*. Whatever exactly these are or turn out to be, they first arose during the 1970s, with a triple claim: to be AI systems that used large bodies of heuristic knowledge, to be AI systems that could be applied, and to be the wave of the future'.[685] MYCIN was, in Newell's view, both the original impetus behind and the principal driver of this transformation. It was, he believed, 'the granddaddy of them all—the one that launched the field'.[686]

But even this new field of AI, once ascendant, would ultimately fall out of favour. By the 1990s, researchers largely abandoned the expert systems approach to AI, returning to statistical approaches that did not require explicit enumeration of expert knowledge. The reasons behind this abandonment were many, but they largely centred on challenges that were already apparent during the field's early MYCIN years. Again and again, developers ran up against the challenge of having to render implicit forms of knowledge explicit, and of having to maintain unruly, ever-expanding knowledge bases. The struggles that plagued this next wave of expert systems engineers would have therefore been familiar to members of the so-called MYCIN gang, who

---

[684] Len Shustek, 'An Interview with Ed Feigenbaum', *Communications of the ACM* 53, no. 6 (2010): 41-45.
[685] Allen Newell, 'Forward' in Bruce G. Buchanan and Edward H. Shortliffe, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (Reading, MA: Addison-Wesley, 1984), xi-xvi, p. xi.
[686] Ibid.

themselves spent countless hours trying to 'extract' and 'crystallise' sometimes implicit and poorly-defined areas of knowledge; and to update, modify, debug, and expand an untamed slurry of IF-THEN rules.

Despite the ultimate fate of MYCIN, and of the larger transformation it wrought, the history of MYCIN brings out tensions, topics, and themes that transcend this particular computerised system, this particular approach to AI, and this particular time period. As computers entered the realm of medical diagnosis and decision making for the first time, matters of trust and transparency emerged as sites of intense moral tension. How could human physicians hope to fully understand a computerised system that took many individuals, working across many years and many disciplines, to develop? How could physicians be expected the place matters of life and death in the hands of faceless algorithms? What kinds of biases might become calcified by an opaque, computerised algorithm? Such questions drove the development of, and responses to, MYCIN, as this section has shown. Almost from the outset, members of the MYCIN team were determined to create a system whose medical reasoning was clear and transparent, subject to the physician's appraisal and, ultimately, rejection. This determination, I have argued, stemmed not just from concerns about transparency but also from concerns about acceptance and authority. The purpose of explanation was to convey *both* information and deference. In developing MYCIN and its explanation capabilities, Shortliffe and his colleagues inscribed certain assumptions about the nature of medical knowledge (explicit, modular)—and about the identity of physicians (atomistic, controlling).

Shortliffe and his colleagues largely pursued *technological* solutions to this problem of trust and transparency—a problem that had moral, social, legal, and epistemological dimensions. Depending on the specific instantiation of the project,

these technological fixes took on different names: with MYCIN it went by 'Subprogram 2'; with ONCOCIN it travelled under 'The Interviewer'. Yet these technological fixes invariably fell short. To deliver IF-THEN rules to physician-users, no matter how sophisticated the user interface, was not to instil a true understanding of MYCIN's reasoning process or its broader domain. In a 1993 essay, Clancey described this as the gap between 'hav[ing] a representation in your pocket' and 'hav[ing] intelligence'. Rules themselves, isolated fragments of medical knowledge, could be difficult to understand in isolation, and in concert they could function in ways that were surprising and non-modular. Clancey elaborated, 'In April 1975, I wrote an essay, "Why is the Tetracycline rule so difficult to understand?" This simple rule presented itself as a puzzle: "If the patient is less than 7 years old, then do not prescribe Tetracycline". I knew what all the words meant, but I couldn't understand why the rule was correct. How could Mycin understand the rule? What did it mean to understand a rule? And if Mycin did understand a rule, how could it be said to be reasoning?'[687] This line of questioning prompted Clancey to reconsider whether exhibition of MYCIN's rules constituted meaningful explanation: 'Somewhat fatuously, I labeled my line-printer listing of rules "What Mycin Knows". More than a decade would pass before I realized that to have a representation in your pocket is not to be intelligent, that "having knowledge" does not mean possessing some *things*. A more accurate label for the rule listing would have been "Our representations of what Drs. Shortliffe, Blum, and Yu know"'.

Clancey's reflections highlight the epistemological barriers to MYCIN's explanation capabilities. But other barriers existed as well. Medical diagnosis is not just about knowledge and information; diagnoses are made within a rich bureaucratic,

[687] Clancey, 'Notes on "Epistemology of a Rule-Based Expert System', p. 197.

moral, social, professional, and economic context, and they perform social and moral, as well as medical, work. As Charles Rosenberg has described, diagnosis 'is a ritual that has always linked doctor and patient, the emotional and the cognitive, and, in doing so, has legitimated physicians' and the medical system's authority while facilitating particular clinical decisions and providing culturally agreed-upon meanings for individual experience'.[688] Inevitably, then, relegating the task of diagnosis to machine was more than an informational challenge. MYCIN's technological capacity to 'explain itself' did little to resolve the profound moral uncertainties of ceding important medical decisions to a machine, nor did it address broader legal and regulatory questions around blame, responsibility, and oversight. Transparent display of MYCIN's rules did not guarantee that certain assumptions and biases were not still 'hidden in the code'.[689]

[688] Charles Rosenberg, 'The Tyranny of Diagnosis: Specific Entities and Individual Experience', *Milbank Quarterly* 80, no. 2 (2002): 237-260, p. 240.
[689] As one anonymous referee would put it to Shortliffe in 1988, '[T]he *major* stumbling blocks (apart from money) have been and will continue to be linguistic, organizational, sociological, ethico-legal, and systematic'. Helen Goldstein to Shortliffe, 20 Jun. 1988, ESP/9.

CONCLUSION AND EPILOGUE

Looking Forward, Looking Back

On 11 April 2018, the U.S. Food and Drug Administration (FDA) announced marketing approval for the device IDx-DR, a software programme developed by an Iowa-based biotech company that detects diabetic retinopathy: 'The U.S. Food and Drug Administration today permitted marketing of the first medical device to use artificial intelligence to detect greater than mild level of the eye disease diabetic retinopathy in adults who have diabetes'.[690] IDx-DR thus became the first FDA-approved fully autonomous diagnostic system in any area of medicine; the software's results do not need secondary approval or interpretation by a human physician. The device uses an AI algorithm to analyse patient retinal images captured by the Topcon NW 400 camera for signs of diabetic retinopathy. By allowing for diagnosis of this eye disease without a visit to the ophthalmologist—on the basis of images that can be captured at primary care offices—the developers hope that the device will save the vision of countless patients through earlier detection of this disease.

The FDA's decision followed the results of a multi-site study of 900 patients known to have diabetes but with no known diagnosis of diabetic retinopathy. The device's performance was impressive. Compared to the gold standard of human grading of diabetic retinopathy, the device diagnosed with a sensitivity of 87.2 percent, a specificity of 90.7 percent and an imageability rate of 96.1 percent, exceeding 'all

[690] 'FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems', FDA News Release, 11 Apr. 2018. Accessed 23 July 2019. Accessible online at: https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye

pre-specified superiority endpoints'.[691] The findings, the clinical trialists concluded, demonstrated 'AI's ability to bring specialty-level diagnostics to primary care settings'.[692]

The U.S. primary care setting hasn't been the only target for AI systems like IDx-DR. Similar technologies are being applied internationally to under-resourced settings. The Aravind Eye Hospital in Madurai, India is currently looking to automated AI systems to ease what it sees as a growing 'bottleneck' problem, with massive numbers of diabetic patients and few qualified ophthalmologists to evaluate them.[693] Partnering with a team of Google AI researchers, the hospital is testing an algorithm—developed jointly by Google and Verily—that can detect diabetic retinopathy. The system, Aravind's chief medical officer argues, will allow physicians 'more time to work closely with patients on treatment and management of their disease while increasing the volume of screenings we can perform'.[694]

This application of AI to ophthalmologic diagnosis, though notably far along, is by no means unique.[695] AI methods are rapidly being developed to help dermatologists identify skin malignancies from skin images, radiologists diagnose pulmonary tuberculosis and other lung diseases from chest radiographs, and pathologists detect

---

[691] For the purposes of calibration against human performance, studies have shown that board-certified ophthalmologists using indirect ophthalmoscopy only reach a sensitivity of between 33 and 73 percent when compared against the same standard.
[692] Michael D. Abràmoff, Philip T. Lavin, Michele Birch, Nilay Shah, and James C. Folk, 'Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices', *npj Digital Medicine* 1 (2018).
[693] Cade Metz, 'A.I. Helps Indian Doctors Save Eyesight', *New York Times* 11 Mar. 2019, B1.
[694] Dr. R. Kim, quoted in Chris Mills Rodrigo, 'Google launches effort in India to screen for leading causes of blindness', *The Hill*, 26 Feb. 2019. Accessed July 23, 2019. Accessible at: https://thehill.com/policy/technology/431591-google-launches-effort-to-screen-for-leading-causes-of-blindness
[695] For an excellent review of recent developments in medical AI, see Kun-Hsing Yu, Andrew L. Beam, and Isaac S. Kohane, 'Artificial Intelligence in Healthcare', *Nature Biomedical Engineering* 2 (2018): 719-731.

prostate cancer from biopsy specimens.[696] Others are applying these methods not just to diagnose but also to predict: to forecast which patients will end up in the ICU, or which chemotherapy recipients will survive after treatment—and for how long.[697] 'Big data', many frequently opine, 'will transform medicine', 'improv[ing] the ability of health professionals to establish a prognosis', 'improv[ing] diagnostic accuracy', perhaps even 'displac[ing] much of the work of radiologists and anatomical pathologists'.[698] With recent strides in AI and big data, some, like the prominent cardiologist Eric Topol, are forecasting—and hailing—the coming 'creative destruction of medicine'.[699]

Rather than simply rehashing this dissertation's arguments and contributions, this conclusion and epilogue instead pulls together and synthesises a number of themes that run through the sections of this dissertation and that bear directly on contemporary efforts to integrate AI and machine learning into medical care. As engineers, physicians, tech companies, and policy makers look forward and try to bring AI tools into the clinic, it is also helpful to look back—to understand that many of the hopes,

---

[696] A. Esteva et al., 'Dermatologist-level classification of skin cancer with deep neural networks', *Nature* 542 (2017): 115–118; P. Lakhani and B. Sundaram, 'Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks', *Radiology* 284 (2017): 574–582; G. Litjens et al., 'Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis', *Science Reports* 6 (2016): 26286.

[697] K.C. Vranas, J.K. Jopling, T.E. Sweeny, M.C. Ramsey, A.S. Milstein, C.G. Slatore, G.J. Escobar, V.X. Liu, 'Identifying Distinct Subgroups of ICU Patients: A Machine Learning Approach', *Critical Care Medicine* 45, no. 10 (2017): 1607-1615. Aymen A. Elfiky, Maximilian J. Pany, Ravi B. Parikh, and Ziad Obermeyer, 'Development and Application of a Machine Learning Approach to Assess Short-term Mortality Risk Among Patients with Cancer Starting Chemotherapy', *JAMA Network Open* 1, no. 3 (2018): e180926.

[698] Ziad Obermeyer and Ezekiel J. Emanuel, 'Predicting the Future—Big Data, Machine Learning, and Clinical Medicine', *New England Journal of Medicine* 375 (2016): 1216-1219.

[699] Eric Topol, *The Creative Destruction of Medicine: How the Digital Revolution Will Create Better Health Care* (New York: Basic Books, 2012).

fears, challenges, and tensions that we face in this so-called era of 'big data' are nothing new.

HOPES

Medicine in the mid-twentieth century faced a series of perceived crises. There was a crisis of information. Medical knowledge, it seemed, was growing at an unmanageable rate and showed no sign of letting up; there was simply too much for the individual physician to know. There was a crisis of manpower. The number of physicians seemed stagnant, even as the U.S. population—and its health needs—climbed steadily upwards; there were simply too many patients for the individual physician to see. There was a crisis of distribution. Whereas metropolitan centres had easy access to health information and services, rural parts of the country seemed increasingly isolated and cut-off; there was simply too much ground for the individual physician to cover.

For early developers of computerised diagnostic systems, the computer offered an alluring technological fix for the pressing social, economic, and distributional challenges of the time. By consolidating and bringing vast amounts of information to the physician's fingertips, computers promised to ease the cognitive load of physicians inundated by growing stores of medical knowledge. By taking over certain rote and routine aspects of clinical care, they promised to increase the physician's clinical capacity—allowing the physician to see more patients, or to focus more on the 'human' aspects of medical care. By linking large medical centres with rural health clinics, they promised to facilitate the distribution of medical knowledge and care to underserved regions.

A half-century later little has changed. Medicine's crises—and the computer's imagined role as a solution—track closely to those during Brodman, Zworykin, and Shortliffe's times. In a 2017 issue of the *New England Journal of Medicine*, Ziad Obermeyer and Thomas Lee argue that medicine has reached a tipping point, with the human mind newly outstripped by the vastness and complexity of medical knowledge: 'The complexity of medicine', they write, 'now exceeds the capacity of the human mind. Computers are the solution…'.[700] Others today see a future with computers taking over the routinised aspects of medical care and giving physicians more time to build relationships with their patients at the human level. In his 2019 book *Deep Medicine*, the cardiologist Eric Topol argues that AI has the potential to reverse the progressive dehumanisation of medicine and restore its 'human side'.[701] Beyond restoring medical services where they already exist, AI, many today argue, also promises to bring medical care to places and patients without existing or adequate care. This historical analysis does not suggest that AI and other technological solutions to such challenges are misplaced or unimportant. But the persistence of these perceived problems—and the persistence of the computer as a perceived solution—suggests that other social and economic factors must be addressed as well.

FEARS

Though an exciting prospect for some, the notion of computers taking over the work of diagnosis also stirred an inner disquietude among many physicians during the mid-twentieth century. These critics saw the computer as an agent of medicine's

[700] Ziad Obermeyer and Thomas H. Lee, 'Lost in Thought—The Limits of the Human Mind and the Future of Medicine', *New England Journal of Medicine* 377 (2017): 1209-1211, p. 1209.
[701] Eric Topol, *Deep Medicine: How Artificial Intelligence Can Make Health Care Human Again* (New York: Basic Books, 2019).

dehumanisation. This was a technology, they warned, that would distance physicians from their patients, denigrate physicians' intuitive skills, hasten the bureaucratisation of medicine, and treat patients as data rather than individuals. Animating many of these critiques were concerns about deskilling and replacement. When computers began making medical decisions, what parts of medicine would remain under the physician's remit?

Recent years have seen the amplification of similar fears. The AI pioneer Geoffrey Hinton has claimed, 'It's quite obvious that we should stop training radiologists'—a view that has received backing from both leading AI engineers like Stanford's Andrew Ng and leading clinicians like UCSF's Robert Wachter.[702] Such claims haven't fallen on deaf ears. One recent study found that fears about displacement of radiologists by AI may be driving many medical trainees away from the field of radiology.[703] With digital enthusiasts prophesying the complete transformation, if not end, of entire fields of medicine—radiology, pathology, dermatology—many physicians and physicians-in-training have expressed concern over how AI will transform, remake, or even eliminate existing jobs and roles.

Beyond fears of deskilling and displacement, others have expressed unease about the computer's potential role in the erosion of medicine's humanistic foundation. These commenters worry that the spread of AI through medicine will replicate many of the problems around the implementation of electronic medical records (EMRs). Now ubiquitous, EMRs have, many physicians believe, disrupted doctor-patient

[702] Quoted in 'Images aren't everything,' *The Economist* 427, no. 9095 (9 Jun. 2018): 15-16, p. 15.
[703] Bo Gong, James P. Nugent, William Guest, William Parker, Paul J. Chang, Faisal Khosa, Savvas Nicolaou, 'Influence of Artificial Intelligence on Canadian Medical Students' Preference for Radiology Specialty: A National Survey Study', *Academic Radiology* 26, no. 4 (2019): 566-577.

relationships, upended professional rituals, increased administrative burdens, and contributed to feelings of burnout. The physician-writer Abraham Verghese has documented what he sees as the EMR-driven dehumanisation of the art of medicine, with physicians' increasingly treating the 'iPatient'—the patient of the EMR—at the expense of the real patient. Verghese writes,

> The patient is still at the center but more as an icon for another entity clothed in binary garments: the "iPatient." Often, emergency room personnel have already scanned, tested, and diagnosed, so that interns meet a fully formed iPatient long before seeing the real patient. The iPatient's blood counts and emanations are tracked and trended like a Dow Jones Index, and pop-up flags remind caregivers to feed or bleed. iPatients are handily discussed (or "card-flipped") in the bunker, while the real patients keep the beds warm and ensure that the folders bearing their names stay alive on the computer.[704]

Developers of AI systems, Verghese argues with colleagues, should pay heed: 'The lessons learned with the EMR should serve as a guide as artificial intelligence and machine learning are developed to help process and creatively use the vast amounts of data being generated in the health care system'.[705] Verghese's admonitions, while prescient, carry forward a line of thinking that has been in the making for more than half a century.

BIASES

The histories of the MDS, HEME, and MYCIN all highlight how data *get made*. None of the actors in this dissertation stumbled upon complete, cleaned, and fully-formed data. Instead, they *created* this data—employing all kinds of assumptions and decisions in the process. The data on which the MDS was based, for example, reflected

---

[704] Abraham Verghese, 'Culture Shock—Patient as Icon, Icon as Patient', *New England Journal of Medicine* 359, no. 26 (2008): 2748-2751, p. 2749.
[705] Abraham Verghese, Nigam H. Shah, and Robert A. Harrington, 'What This Computer Needs Is a Physician: Humanism and Artificial Intelligence', *JAMA* 319, no. 1 (2018): 19-20.

important decisions that Brodman had to make in developing the CMI: his choices about which questions to include, and his decisions about which population of patients on which to test the questionnaire. The data behind certain versions of the haematology programme also rested on subjective human judgements. When, for example, the team looked to the medical literature for data on which to base their programme, they discovered a remarkable absence of clarity and uniformity. Different clinicians and researchers reported medical data differently; they held to varying levels of quality, used different disease names, and fell at different places along the qualitative-quantitative spectrum. Making sense of such data—extracting it, processing it, combining it—was no simple matter. It required the team to make subjective, ad-hoc decisions about which studies to include and exclude, how to resolve contradictory reports, and whether and how to quantify qualitative data. The team turned to the medical literature to eliminate the need for human judgement—to do away with physician-generated estimates of significance in favour of 'objective' significance data. Yet the need for subjective judgement remained. In the case of MYCIN, data was made and constantly re-made: the team was engaged in a process of relentless maintenance, always adding, subtracting, and tweaking the system's IF-THEN rules.

Attention to the contingencies and circumstances around the creation of data highlights the challenges of error and bias in computerised systems. Although many mid-twentieth-century proponents of computerisation viewed computerised systems as inherently objective, others were alive to possible problems of error and bias. These actors warned of the possibility of inaccurate reference data; they adopted the mantra, 'garbage in, garbage out'.

AI in healthcare today depends on the availability of massive, high-quality datasets. Yet insufficient attention is often paid to how this data gets made. What kinds

of biases might this data reflect or entrench? What kinds of decisions have gone into the collection and cleaning of various datasets? A historical perspective foregrounds the messy and monotonous, contingent and consequential work of data production.

Medical data is extraordinarily complex. For it to be useful and meaningful, some kind of order must be imposed on it. In the process of imposing that order, however, decisions and judgements must be made, and biases and errors can be introduced. The use of big data in medicine has not erased the necessity for clinical and human judgement but rather defined new areas for its application. What assumptions, biases, and errors might a programme encode and propagate? To what degree does a given programme, and the data on which it was produced, apply to a given patient or patient population? These were questions that physicians first grappled with in the mid-twentieth century, and they are questions that clinicians will continue to grapple with well into the twenty-first century.

These are not just hypothetical issues. Medical studies often suffer from relatively small numbers of racial and ethnic minorities. As a result, key sources of medical data often do not adequately represent all groups within society equally. AI systems based on such data may produce medical decisions suited for white patients, but not black patients, for example. Similar scenarios have already been borne out by existing technologies. A machine learning algorithm developed to predict the risk of criminal recidivism has displayed glaring racial biases.[706] Facial recognition software has performed worse on non-white faces because this software was developed with unrepresentative data.[707] A historical perspective helps alert us to the real possibility of

[706] Julia Dressel and Hany Farid, 'The accuracy, fairness, and limits of predicting recidivism', *Science Advances* 4, no. 1 (2018): eaao5580.
[707] Joy Buolamwini and Timnit Gebru, 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', *Proceedings of Machine Learning Research* 81 (2018): 1-15.

AI calcifying bias, and exacerbating disparities, in our health care system. To the extent that history helps us understand where data comes from—how data 'gets made'—it also represents an important tool as we work to avert and remedy the problems of bias.

STANDARDS

It is a commonplace of evidence-based medicine to speak of 'the gold standard'—of the randomised controlled trial as the gold standard of evidence, of *this* or *that* procedure as the gold standard of practice. Yet the history of computerising diagnosis demonstrates how elusive gold standards of diagnosis can be. After early developers created system prototypes, they frequently wanted to test these prototypes. How well did they perform compared to human physicians? Immediately, though, the developers realised that often in medicine gold standards of performance are difficult to identity and define. All too often there is no Supreme Court in medical decision making. Instead, people like Brodman, Shortliffe, and Engle encountered a world of medicine where experts disagree and vary in their judgements. In response to these realities, this dissertation's actors pursued methodological innovations and clever research designs, for example, comparing the computer's performance against varying gradations of expertise and using 'super-experts' to evaluate the evaluators. Yet the fundamental fact remained that whether the computer was 'correct' could not be known with perfect certainty.

Today, standards continue to present difficulties for those working to computerise medical decision making. In machine learning systems, developers must use a training data set that defines 'ground truth' for the algorithm: the machine learning algorithm uses the training set to mathematically define borders between cases on the basis of already-known classifications in the training data. In medicine,

confidently defining 'ground truth'—arriving at these already-known classifications—is no simple matter.[708] Because data about disease diagnosis can be messy and inaccurate, based in large part on subjective human judgements, the training data for ground truth can be off. Almost always, then, ground truth in medicine is fuzzy. Moving out from diagnosis in particular to decision making more broadly, medical decisions are often not matters of black-and-white fact but rather matters of morals and values: the 'right' treatment, for instance, depends not just on the results of randomised controlled trials but also on the goals and circumstances of the patient. If, how, and whether computers should be applied in such situations remain open and hotly contested questions.

OPACITIES

Through the mid-twentieth century, pioneers of computerised diagnosis debated different approaches to digitising medical decision making. Some looked inward, trying to define, explicate, and ultimately computerise how human physicians go about the process of making a diagnosis. Others looked outward—to new, non-human modes of reasoning. To this crowd, mimicry made little sense: not only are the strengths of the computer different from those of the human but humans may make medical decisions in suboptimal ways. Computerise diagnosis, they proposed, in a way that maximised the relative strengths of the machine and minimised the relative weakness of the human. Central to the debate over the relative merits of these different strategies was the question of interpretability. Many argued that by replicating human reasoning processes, the computer would generate decisions that humans could understand,

[708] Sendhil Mullainathan and Ziad Obermeyer, 'Does Machine Learning Automate Moral Hazard and Error?', *American Economic Review* 107, no. 5 (2017): 476-480.

evaluate, affirm, or override. Other, more statistically oriented modes of reasoning, by contrast, risked looking opaque and foreign—uninterpretable. In the former scenario, some believed, physicians would be more willing to accept and adopt computerised systems, as these technologies would be more deferential to their authority and autonomy.

Much work went into laying bare the computer's reasoning for the physician, as the MYCIN case makes clear. Shortliffe and his colleagues developed sophisticated technologies that allowed MYCIN to 'explain itself'. Their hope was that by delivering to the physician the decision rules that MYCIN had used in its decision making process physicians would find the system more acceptable. Yet even with this technological capability, problems and concerns remained. Some questioned whether the display of IF-THEN rules really constituted meaningful explanation. Others continued to worry about possible biases 'hidden in the code' of MYCIN and how these biases might undermine the physician's personal responsibility to the individual patient. These technological solutions, moreover, did not resolve open legal and regulatory questions concerning liability.

Interpretability has only become more of an ethical, legal, and professional flashpoint in recent years. Because machine learning algorithms create their own decision patterns and criteria on the basis of vast quantities of data, it is often impossible for humans to understand the processes by which AI systems arrive at their decisions. Motivated by novel digital challenges like those of the 'black box', the EU rolled out a General Data Protection Regulation (GDPR) in 2018 that required AI systems to be interpretable. In medicine—a domain where machine learning may be applied not to matters of advertising revenue and click-through rates but potentially to matters of life and death—the problem of interpretability is particularly salient. Many

commentators believe that turning medical decisions over to uninterpretable machines runs counter to the physician's moral and professional duties. Others point out that other core medical tenets—such as those of patient disclosure and safety—may also be challenged by AI systems that are making medical decisions in the dark.[709]

Not all, however, are so worried. Some have countered that concerns about interpretability are overblown—that the validity of computer-generated decisions should be held up as a higher social value.[710] Human intelligence, they point out, is often unintelligible and inexplainable, and society already tolerates all kinds of black boxes in medicine, such as the mechanisms of all kinds of drugs.[711]

While many researchers are pursuing promising technological solutions to the problem of interpretability—rolling out natural language processing technologies, developing AI systems to interpret the results of the original AI system, and defining the new research area of eXplainable AI (XAI)—history suggests that the challenges introduced by opaque AI systems will not be solved through technological solutions alone. Interpretability is a value with ethical, legal, and professional valences. For example, if making an AI system interpretable comes at the expense of its predictive accuracy (as is often the case), what should be done? The problems around interpretability can only be addressed through the collective moral, legal, and

---

[709] Effy Vayena, Alessandro Blasimme, I. Glenn Cohen, 'Machine Learning in Medicine: Addressing Ethical Challenges', *PLOS Medicine* 15, no. 11 (2018): e1002689. On the legal issues introduced by black box algorithms, see, e.g., W. N. Price, 'Regulating Black-box Medicine', *Michigan Law Review* 116, no. 1 (2017): 421–474.

[710] See, e.g., Elizabeth Holm, 'In Defense of the Black Box', *Science* 364, no. 6435 (2019): 26-27.

[711] See, e.g., Alex John London, 'Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability', *The Hastings Center Report* 49, no. 1 (2019): 15-21.

regulatory work of articulating—and implementing—societal standards of transparency.

IDENTITIES

In early efforts to computerise medical diagnosis, the computer functioned—to borrow a phrase—as 'an engine, not a camera'.[712] It didn't just capture and implement preexisting notions of medical reality. It also generated new notions about what medicine was and what it could be. It opened up a space where physicians and engineers began to rethink and redefine fundamental categories and roles. The computer's first forays into medical diagnosis prompted people like Brodman, Zworykin, Engle, and Shortliffe to reflect on the physician's identity—to identify and define the aspects of medical work and medical thought that belonged to human physicians and those that belonged to machine. Always, ideas about the computer's 'appropriate' role in the clinical environment reflected, reinforced, and displaced certain ideas about the nature and identity of humans.

Early computerised diagnostic programmes were also mutually constitutive with respect to disease identities. The MDS, for example, both gave rise to and depended upon conceptions of disease that translated the 'symptom complex' into the language of numbers, statistics, and thresholds. The haematology team similarly generated new ways of defining haematological disease through their engagements with new digital media.

Questions about how computers may reconfigure disease, professional, and patient identities remain relevant today. As computers take on new clinical tasks, the

[712] Donald MacKenzie, *An Engine, Not a Camera: How Financial Models Shape Markets* (Cambridge, MA: MIT Press, 2006).

kinds of work deemed 'appropriately human' may shift. Already computing technologies have transformed how certain medical personnel are seen—and how they see themselves. In the 1990s, radiological reading rooms around the country adopted a technology called picture archiving and communication system (PACS). A medical imaging capability, PACS provides radiologists and other clinicians ready, computerised access to radiological images and their associated interpretations and reports. Following the implementation of PACS, the radiological reading room—once a place bustling activity, a place where physicians of all kinds would come to discuss a case—became removed and isolated from the rest of the hospital.[713] Because of PACS, the physical objects that drew clinicians of all stripes to this space—slides or film—were no more.

This PACS-mediated transformation in patterns of interaction at the hospital has influenced the radiologist's identity. Many radiologists have expressed concern about new limitations in their ability to 'communicate clinical recommendations, relevant research, or alternate diagnostic options with the patient care team' following the adoption of PACS.[714] They have also felt increasingly disconnected. One senior radiologist opined,

> We [radiologists] knew all the clinicians intimately before. And then with PACS, this intimacy disappeared. Before [PACS], I knew the face, name, wife's name, and kids' name of all the clinicians, but now I don't know who you are if you joined the medical staff after we got PACS. Now we're operating in a void because there's no history of the patient on the written image requests. Before, when a clinician showed up, I could ask them and find out what's really going on with the patient.[715]

---

[713] Allison A. Tillack and Richard S. Breiman, 'Renegotiating Expertise: An Examination of PACS and the Challenges to Radiology Using a Medical Anthropologic Approach', *Journal of the American College of Radiology* 9, no. 1 (2012): 64-68.
[714] Ibid., p. 66.
[715] Ibid., p. 67.

Computing technologies can also shape the identities of patients and their diseases. How individuals experience disease today differs wildly from how individuals experienced disease in the past. More than ever before, our experience and understanding of disease is mediated through numbers. Though a patient may not subjectively feel sick, he might interact with the medical system as a 'pre-diabetic' because his fasting blood sugar level falls between 100 and 125 mg/dL. Though a patient may not actually have breast cancer, she might make important health and life decisions knowing that, as a BRCA1 mutation carrier, she has a roughly 65 per cent chance of developing breast cancer by the age of 70. Patients such as these experience health and disease not through any extant symptoms but through numbers—and through regimens of medical surveillance and management designed to minimise the health risks associated with these numbers.[716] Many patients today stand before what the anthropologist S. Lochlann Jain evocatively calls 'the firing squad of statistics'.[717]

In addition to the experience of patienthood, the act of diagnosis is also an increasingly numerical enterprise. For most of history, patients' personal feelings and reports of their own health and body formed the basis of disease and diagnosis; in one way or another, a patient diagnosed with 'dis-ease' really did subjectively lack ease. This is not always so for patients today. Whether a patient has a given disease is often a statistical question. Physicians and patients today inhabit a world in which medicine and statistics have converged—a world of 'risky medicine'.[718]

---

[716] See, e.g., Jeremy Greene, *Prescribing by Numbers.*
[717] S. Lochlann Jain, *Malignant: How Cancer Becomes Us* (Berkeley: University of California Press, 2013), p. 27.
[718] Robert Aronowitz, *Risky Medicine: Our Quest to Cure Fear and Uncertainty* (Chicago: University of Chicago Press, 2015); Robert Aronowitz, 'The Converged Experience of Disease and Risk', *Milbank Quarterly* 87, no. 2 (2009): 417-442.

This dissertation prompts consideration of the ways in which computing technologies may be shaping this modern style of defining and experiencing disease. A number of pathways are readily apparent: these technologies permit the complex computations and administer the large databases that underlay a growing web of epidemiological knowledge; they, through automated health screenings and electronic health records, form much of the basis for our contemporary screening apparatus; and they have facilitated the intensification of diagnostic testing through phenomena like the automated processing of tests.

Such possibilities were on the front of the minds of people like Zworykin. In 1967, for example, one of his associates wrote to Zworykin about their shared interest in 'how computer data processing [could be] used to obtain an "early warning system"'—that is, a way of identifying and monitoring 'at risk' patients.[719] Throughout his speeches and writings, Zworykin often spoke of computerised diagnosis and large scale medical monitoring in the same breath. Zworykin foresaw a future in which computer centres storing medical records would collect 'tremendous amounts of clinical data' and 'provide the basis for medical statistical studies on an unparalleled scale'.[720] Such information, he suggested, would allow physicians, researchers, and public health officials to make 'many elusive correlations' among sundry clinical variables.[721] As a result, he ventured that doctors—and perhaps even the wider public—would soon understand the health risks associated with things like 'prenatal

[719] Carl Berkley to Vladimir Zworykin, 22, 1967.
[720] Vladimir Zworykin and Carl Berkley, 'Applications of Medical Data Processing', Paper No. VMe-3-57, VZP/77/69; Robert Haavind, 'Needed: A Medical Technology Institute', 26 Apr. 1965, VZP/82/15.
[721] Zworykin and Berkley, 'Applications of Medical Data Processing'.

radiologic examination' or 'food additives such as antimycotics in bread'.[722] In many respects, the world Zworykin imagined is the world we live in today.

CONTINUITIES

Computerisation is often synonymous with change. It makes things faster and more efficient, alters fundamental workflows, promotes a world of information and interconnection. This dissertation has tracked important changes attendant to the computer's introduction into medical diagnosis—changes in how diseases are defined, how medical information is created and handled, how professional roles and identities are constructed. Yet this dissertation has also traced important continuities that link our current digital world dominated by computers to our former analogue worlds dominated by pencil and paper. Many of the earliest efforts to computerise medical diagnosis evolved directly from paper-based methods and tools. The haematology programme had its roots in Firmin Nash's analogue 'Logoscope' as well as Lipkin and Hardy's punched card method of making haematologic diagnoses. The computer-based Medical Data Screen would have never come into being had the paper-based Cornell Medical Index not existed before it.

The pattern here is one of path dependency and constraint. Decisions made with respect to analogue tools and technologies often carried over into their digital descendants. The fact that Lipkin and Hardy's analogue device made diagnoses by way of 'matching' informed and motivated the matching algorithm that Zworykin and his colleagues ultimately chose to implement as they worked to computerise haematologic diagnosis. Similarly, the nature of the CMI questionnaire affected the functioning of the

[722] Ibid.

computerise MDS method—from the questions it included to the syndromal clusters of symptoms it recognised.

Understanding these continuities has important implications for imagining and implementing digital medicine today. Discussions about the impact of computers in fields such as medicine are not confined to the academic historical literature. Physicians, policy makers, social critics, and others frequently extol the virtues, or bemoan the harms, of computerised medicine. The prominent cardiologist Eric Topol, for example, has hailed the arrival of what he has called 'the creative destruction of medicine', with digital technologies facilitating new modes of medical research and reconfiguring the relationship between doctor and patient.[723] Thanks to computerisation, some suggest, the medicine of today is (unlike the medicine of yesterday) a 'data-driven science'. While commentators may disagree about the appeal of medicine's digital transformation, few quibble about its magnitude: when medicine moved from the analogue world of paper charts to the digital world of EMRs, everything changed. This history instructs us not so much to discard these claims about medicine's digital transformation: there is no question that the introduction of computers into medicine has changed medicine in a number of important, even fundamental, ways. Rather, the cases presented here invites us to engage with these well-worn claims with more thought, scrutiny, and scepticism. This history offers a sobering countermelody to a familiar tune of sharp discontinuity and 'creative destruction'. As narratives of discontinuity and disruption dominate headlines, this story helps us to see the continuities that link our present to the past, to engage with the

[723] Eric Topol, *The Creative Destruction of Medicine: How the Digital Revolution Will Create Better Health Care* (New York: Basic Books, 2012); Eric Topol, *The Patient Will See You Now: The Future of Medicine Is in Your Hands* (New York: Basic Books, 2015).

promises of computerised medicine more critically, and to appreciate that many of the questions and challenges that medicine faces today are not entirely new.

From a practical standpoint, even in our increasingly digital clinical environments, older paper data practices have not only persisted to varying degrees but also shaped how existing digital infrastructures are imagined, built, and used.[724] EMRs have fallen short of their promise and potential, with physicians now spending increasing amounts of time on administrative tasks, wading through excessive 'alerts', and battling climbing levels of burnout. These shortcomings reflect some of the issues of contingency, path dependency, and constraint that this dissertation has highlighted. By and large, EMRs were not developed as a part of the total hospital infrastructure but rather in a fragmented manner, geared more toward billing and administration than toward patient care. Attending to, and grappling with, the legacies of paper technologies and epistemologies in our current digital environment is thus an important step on the way to reimagining and better implementing EMRs today. Modern medicine's empire of machines came into being upon an infrastructure of paper.

COMPLEXITIES

'The making of correct diagnostic interpretations of symptoms', Brodman confidently declared in 1959, 'can be a process in all aspects logical and so completely defined that it can be carried out by a machine'.[725] Yet behind the scenes Brodman struggled with colleagues to work out lingering kinks and imperfections in their

[724] R. J. Baron, 'Quality Improvement with an Electronic Health Record: Achievable, but Not Automatic', *Annals of Internal Medicine* 147 (2007): 548–552; and Katherine Choi, Yevgeniy Gitelman, and David A. Asch, 'Subscribing to Your Patients—Reimagining the Future of Electronic Health Records', *NEJM* 378 (2018): 1960–1962.
[725] Keeve Brodman, A. J. van Woerkom, Albert J. Erdmann, Jr., Leo S. Goldstein, 'Interpretation of symptoms with a data-processing machine', *AMA Archives of Internal Medicine* 103 (1959):776-782.

computerised method—and to open up the 'black box' of how human physicians actually think. Zworykin's group chose to create a diagnostic programme in the field of haematology because of the field's supposed objective, quantitative, and straightforward nature. Yet they quickly found that the field was rich with complexity—and subjectivity. Shortliffe and his colleagues worked to codify much of the field of infectious disease into hundreds of IF-THEN rules. Yet they soon discovered that certain aspects of medical knowledge may be 'tacit'—and that even the non-tacit aspects of knowledge were subject to considerable debate, disagreement, and contestation among medical experts.

This has been a history of failure. The actors followed in this dissertation repeatedly came up against—and underestimated—the complexity of medicine and medical decision making. This complexity came from many places: the idiosyncrasy of the individual patient, the heterogeneity of medical data, health and health care's moral overlay, the opacity of the human mind, the social construction of many of medicine's classifications. While today's era of 'big data' involves very different technologies, capabilities, and actors than the 'big data' moment of the mid-to-late twentieth century, these fundamental complexities remain.

* * *

In Brodman and Zworykin's time, conversations about the role of computers in medical diagnosis found their way into mainstream fora for the first time in history. Either could have picked up the *New York Times* on 8 February 1958 to read the headline 'Diagnosis by Computer Envisaged for Patients'.[726] In this way, their time is not unlike our own. Stories about the possible uses of computers in medical diagnosis

[726] 'Diagnosis by Computer Envisaged for Patients', *New York Times*, 8 Feb. 1958, p. 6.

swarm both medical journals and lay publications.[727] The questions that motivate these accounts tend to be questions of accuracy. To what extent does the computer 'get it right'? What features of medical practice and medical knowledge limit their accuracy? What can be done—from engineering, behavioural, or health systems standpoints—to improve their effective use?

I do not mean to suggest that these are unimportant questions. It is vitally important to understand how to integrate computerised and machine-learning methods into the social world of medicine in ways that promote effective and humane medical care. But it is also important to recognise that computers, in diagnosing disease, do not just identify predefined disease entities. They give shape to what disease is, and how we think about it. They construct new professional roles and identities. They create profound moral tensions around values like precision, trust, and transparency. They eliminate some human biases, while at the same time replicating—and extending—others.

This history offers no clear path forward. It does, however, widen our perspective; it raises pressing questions, problems, and challenges that must be addressed as we work to harness AI for the benefit of the medical profession and its patients.

[727] See, e.g., Steve Lohr, 'IBM, in It for the Long Haul, Is Betting Big on Watson', *New York Times*, 17 Oct. 2016, p. B1; Siddhartha Mukherjee, 'The Algorithm Will See You Now', *The New Yorker*, 3 Apr. 2017; Megan Molteni, 'Want a Diagnosis Tomorrow, Not Next Year? Turn to AI', *Wired*, 10 Aug. 2017; Jeremy Hsu, 'Can a Crowdsourced AI Medical Diagnosis App Outperform Your Doctor?', *Scientific American*, 11 Aug. 2017.

FIGURES

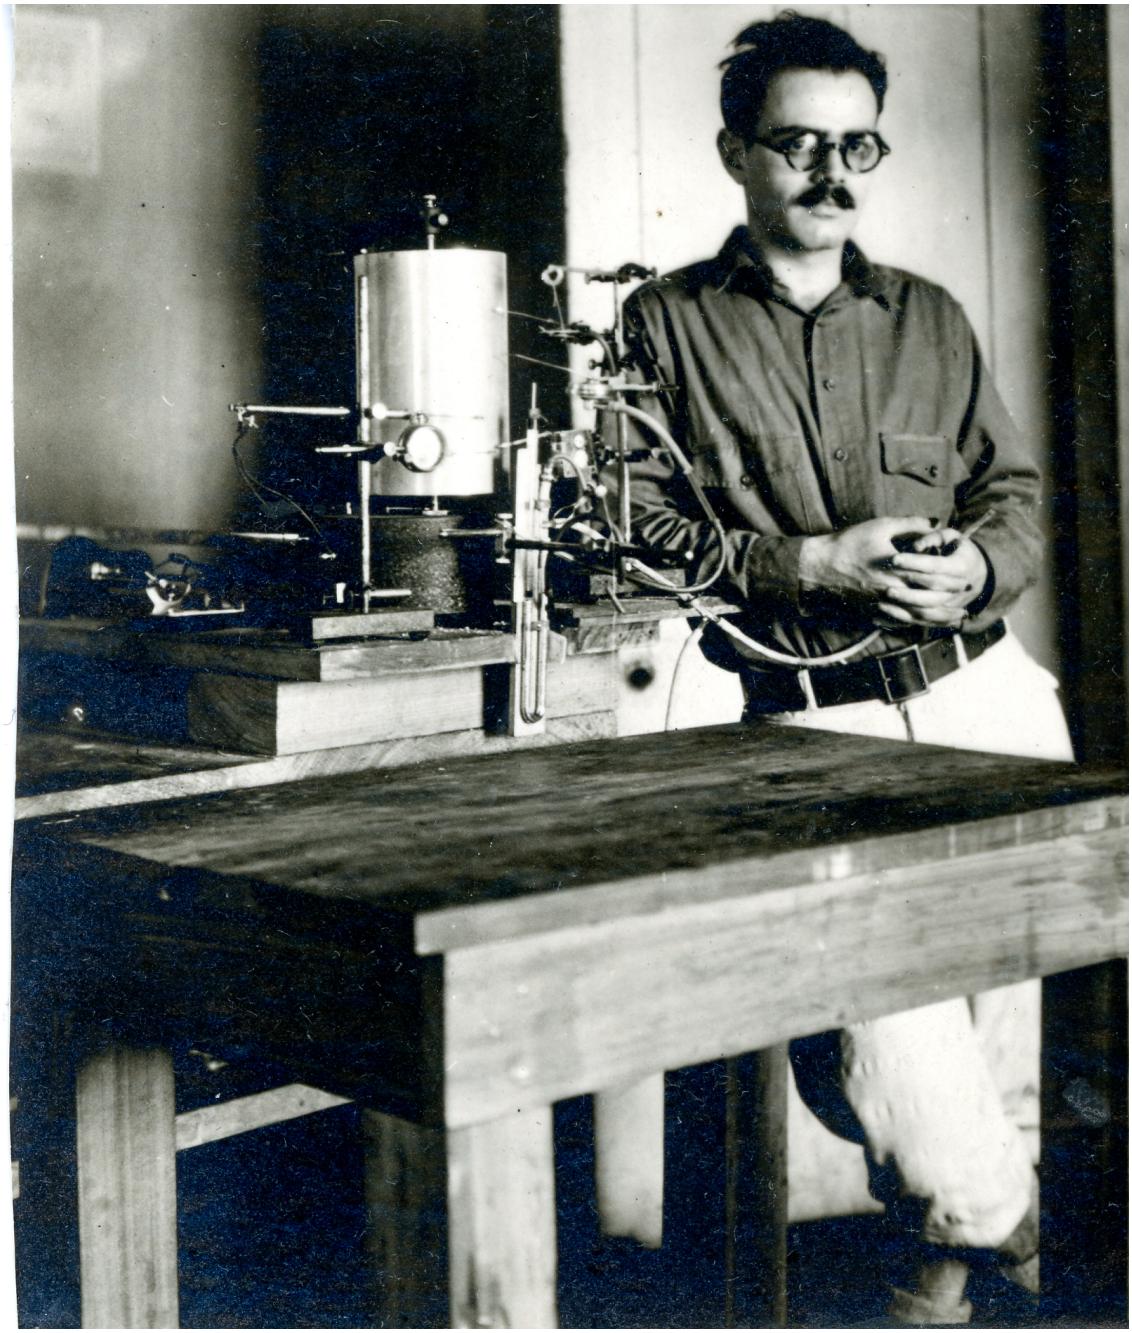

Figure 1: A young Keeve Brodman, 1927. Medical Center Archives of New York-Presbyterian/Weill Cornell.



Figure 2: The New York Hospital-Cornell Medical Center shortly after it opened in 1932. Medical Center Archives of New York-Presbyterian/Weill Cornell.

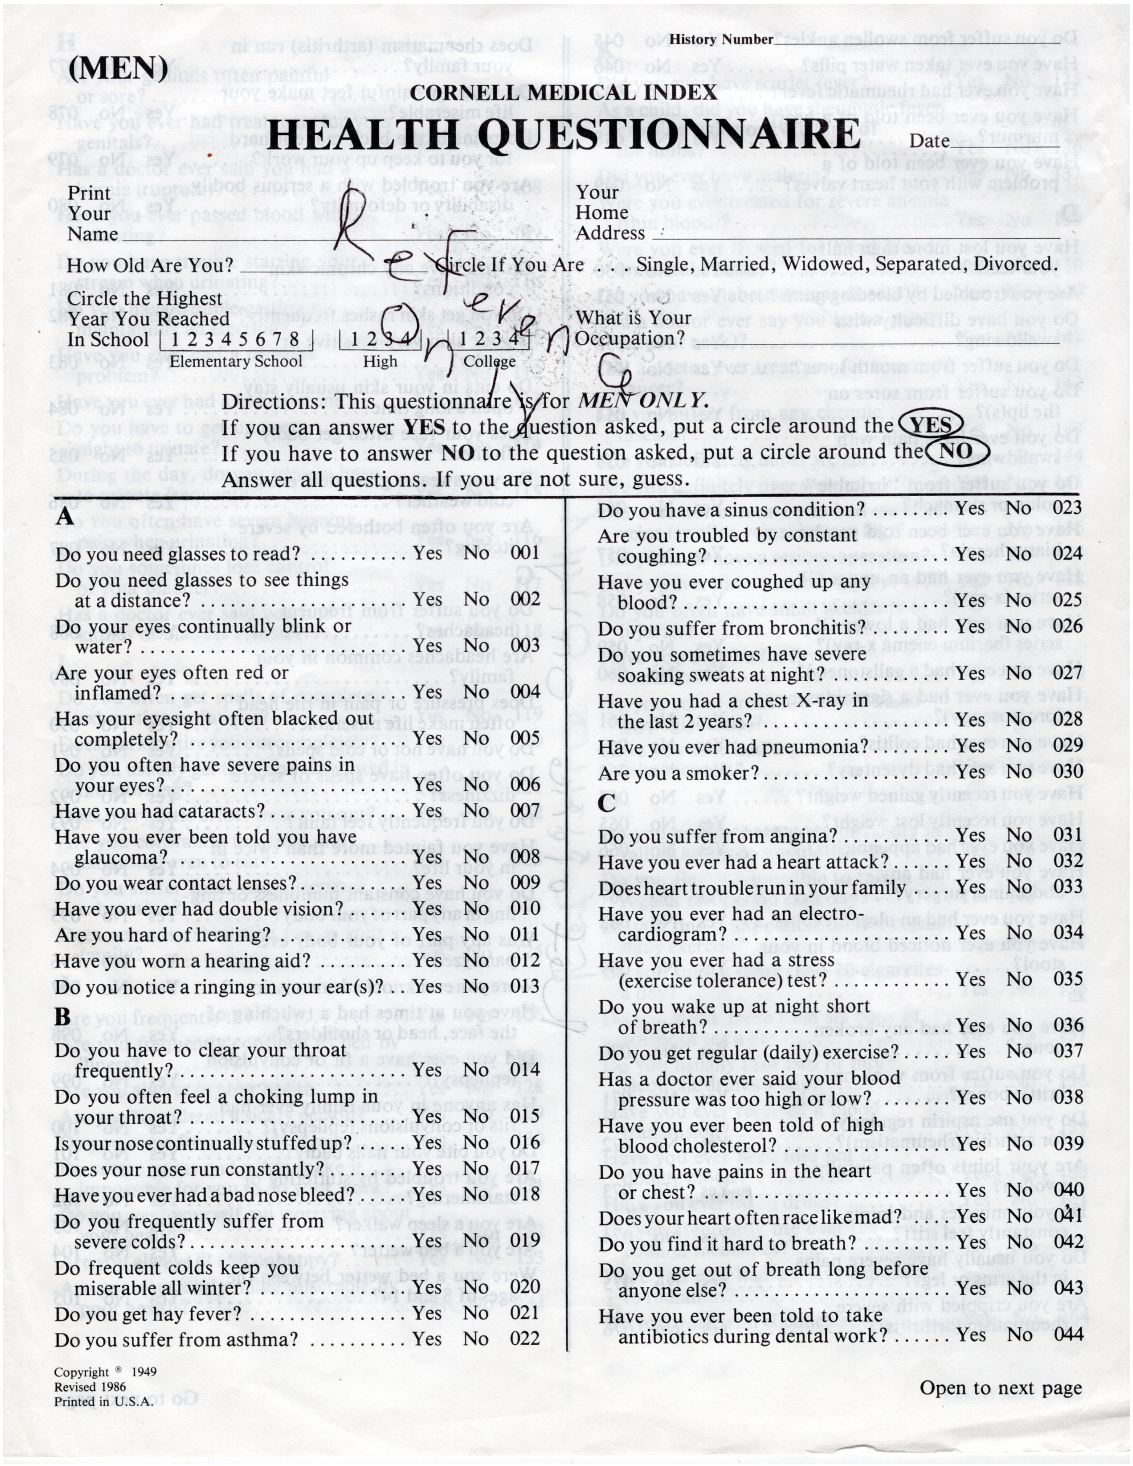

History Number_____

(MEN)

CORNELL MEDICAL INDEX

HEALTH QUESTIONNAIRE

Date_____

Print Your Name _____ Your Home Address _____

How Old Are You? _____ Circle If You Are . . . Single, Married, Widowed, Separated, Divorced.

Circle the Highest Year You Reached In School | 1 2 3 4 5 6 7 8 | Elementary School | 1 2 3 4 | High | 1 2 3 4 | College | What is Your Occupation? _____

Directions: This questionnaire is for *MEN ONLY.*
If you can answer **YES** to the question asked, put a circle around the YES
If you have to answer **NO** to the question asked, put a circle around the NO
Answer all questions. If you are not sure, guess.

A

Do you need glasses to read?	Yes	No	001
Do you need glasses to see things at a distance?	Yes	No	002
Do your eyes continually blink or water?	Yes	No	003
Are your eyes often red or inflamed?	Yes	No	004
Has your eyesight often blacked out completely?	Yes	No	005
Do you often have severe pains in your eyes?	Yes	No	006
Have you had cataracts?	Yes	No	007
Have you ever been told you have glaucoma?	Yes	No	008
Do you wear contact lenses?	Yes	No	009
Have you ever had double vision?	Yes	No	010
Are you hard of hearing?	Yes	No	011
Have you worn a hearing aid?	Yes	No	012
Do you notice a ringing in your ear(s)?	Yes	No	013

B

Do you have to clear your throat frequently?	Yes	No	014
Do you often feel a choking lump in your throat?	Yes	No	015
Is your nose continually stuffed up?	Yes	No	016
Does your nose run constantly?	Yes	No	017
Have you ever had a bad nose bleed?	Yes	No	018
Do you frequently suffer from severe colds?	Yes	No	019
Do frequent colds keep you miserable all winter?	Yes	No	020
Do you get hay fever?	Yes	No	021
Do you suffer from asthma?	Yes	No	022
Do you have a sinus condition?	Yes	No	023
Are you troubled by constant coughing?	Yes	No	024
Have you ever coughed up any blood?	Yes	No	025
Do you suffer from bronchitis?	Yes	No	026
Do you sometimes have severe soaking sweats at night?	Yes	No	027
Have you had a chest X-ray in the last 2 years?	Yes	No	028
Have you ever had pneumonia?	Yes	No	029
Are you a smoker?	Yes	No	030

C

Do you suffer from angina?	Yes	No	031
Have you ever had a heart attack?	Yes	No	032
Does heart trouble run in your family	Yes	No	033
Have you ever had an electro-cardiogram?	Yes	No	034
Have you ever had a stress (exercise tolerance) test?	Yes	No	035
Do you wake up at night short of breath?	Yes	No	036
Do you get regular (daily) exercise?	Yes	No	037
Has a doctor ever said your blood pressure was too high or low?	Yes	No	038
Have you ever been told of high blood cholesterol?	Yes	No	039
Do you have pains in the heart or chest?	Yes	No	040
Does your heart often race like mad?	Yes	No	041
Do you find it hard to breath?	Yes	No	042
Do you get out of breath long before anyone else?	Yes	No	043
Have you ever been told to take antibiotics during dental work?	Yes	No	044

Copyright ® 1949
Revised 1986
Printed in U.S.A.

Open to next page

Figure 3: After its original publication in 1949, the CMI was updated and revised a number of times. The final version of the CMI, printed here, retains the aesthetic character and much of the content of the original questionnaire. On the recommendation of a committee appointed by Cornell's Chairs of Medicine, Neurology, and Psychiatry, the CMI was discontinued beginning in 1990. Today the questionnaire is only available for historical ('reference only') purposes. Medical Center Archives of New York-Presbyterian/Weill Cornell.

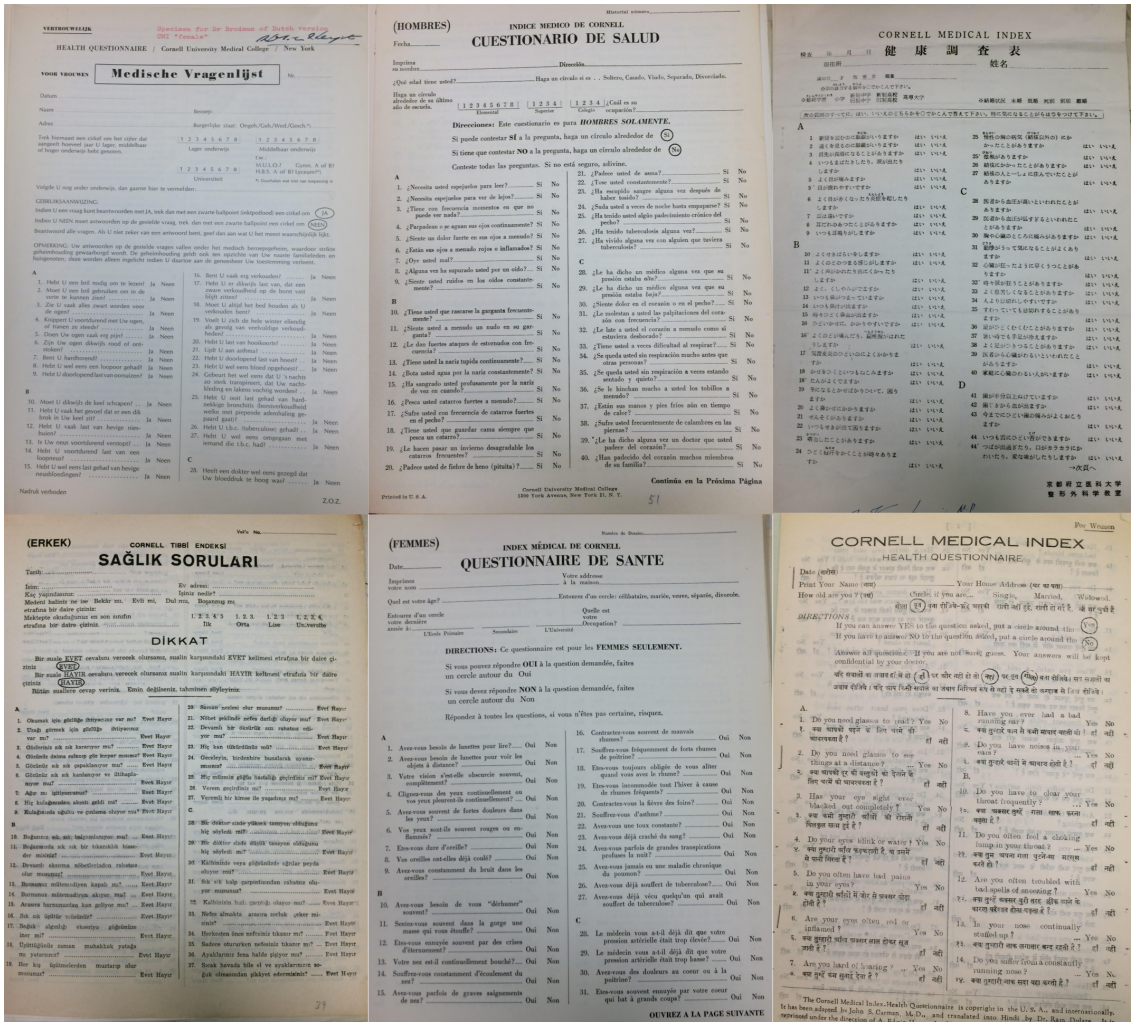

Figure 4: Various translations of the CMI brought the tool to new populations and parts of the world. Medical Center Archives of New York-Presbyterian/Weill Cornell.



Figure 5: IBM 704 Electronic Data Processing System. Computer History Museum.

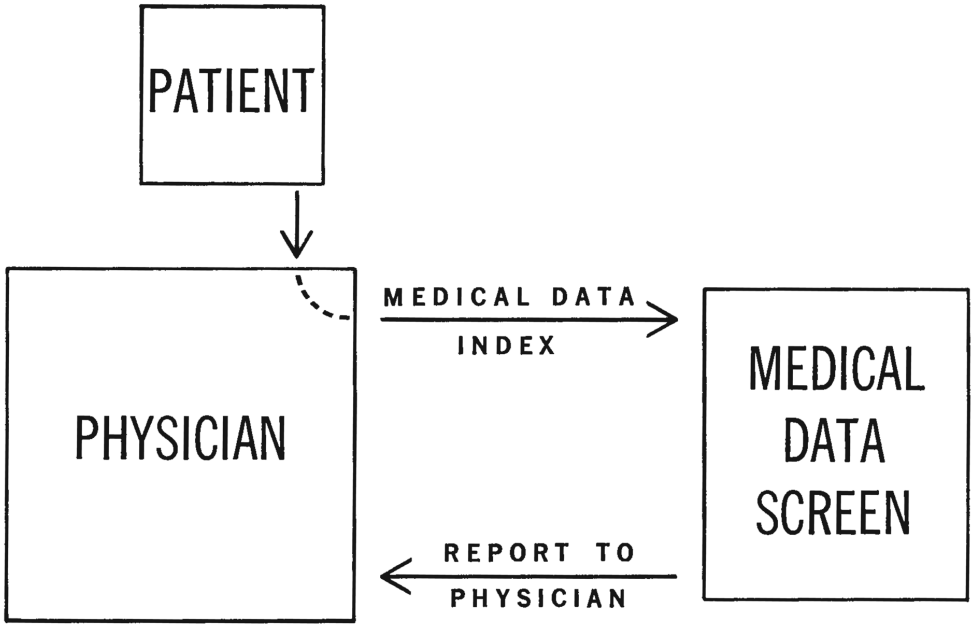

Figure 6: A schematic of the Medical Data Screen's imagined implementation. Medical Center Archives of New York-Presbyterian/Weill Cornell.

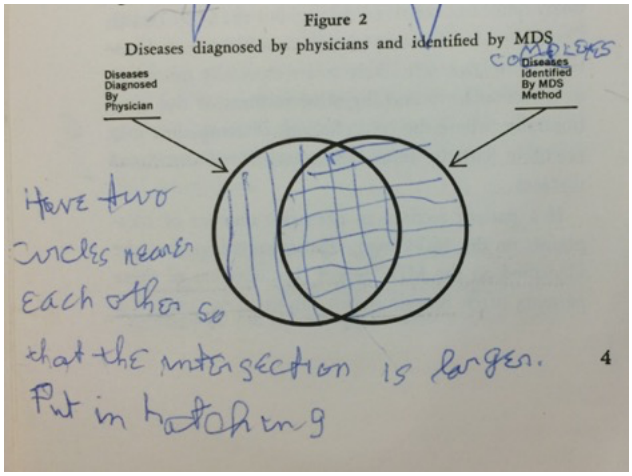

Figure 7: Hand revisions of the Manual made by Brodman reflect the uncertainties around the differences between the category of 'disease' and that of 'symptom complex'. Medical Center Archives of New York-Presbyterian/Weill Cornell.

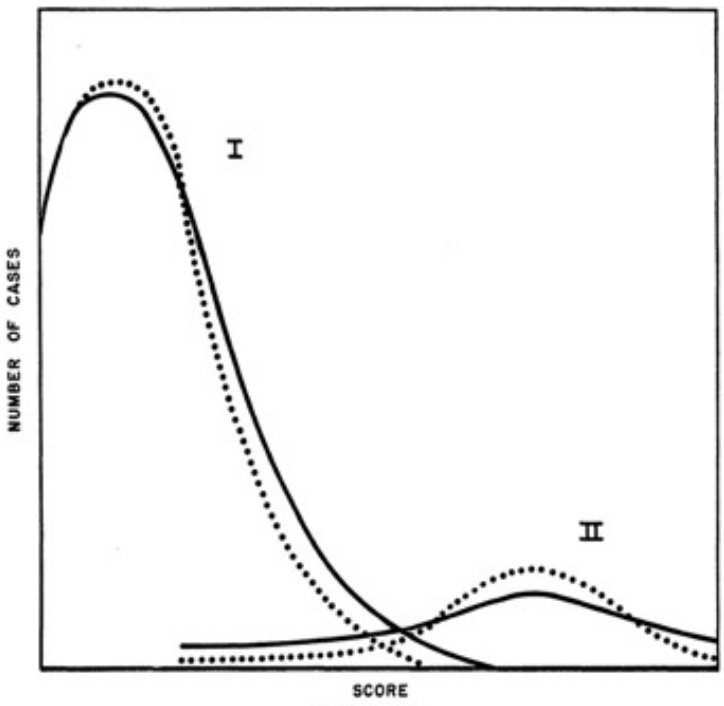

FIGURE 1

DISTRIBUTION OF SCORES FOR CATEGORY AND TOTAL SAMPLE.

Figure 8: The translation of symptoms into the language of numbers meant that individual patient scores could be compared against a normal distribution of scores in the population. Population statistics thus became the reference point for determining health and disease. Adrianus van Woerkom and Keeve Brodman, 'Statistics for a Diagnostic Model', *Biometrics* 17 (1961): 299-318.



Figure 9: Undated aerial photograph of the Radio Corporation of America (RCA) facilities in Camden, New Jersey. Hagley Museum & Library.

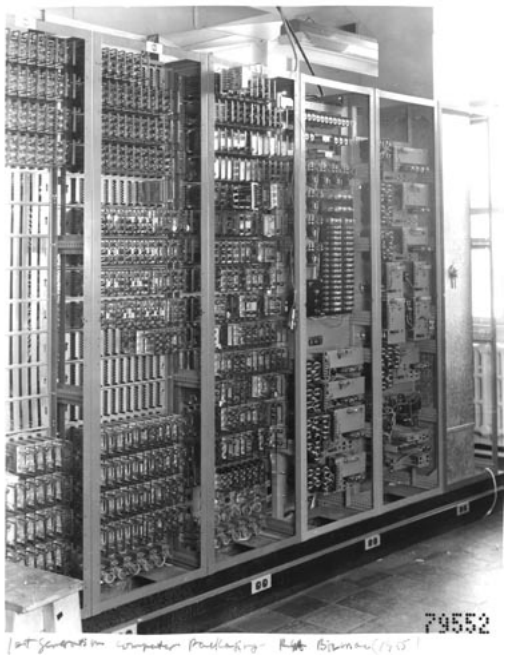

Figure 10: RCA Bizmac frame, 1955. Computer History Museum.



Figure 11: Robert Ledley, pictured with an IBM 360 computer. Creative Commons.



Figure 12: Inventor and television pioneer Vladimir Zworykin (right). Hagley Museum & Library.

TABLE 2.—*Distribution of Information to Marginal Punched Cards: History*

Disease	Code* 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29
Pernicious anemia	+	+	+	+			+	+	+		+	+	+	+				+	+	+	+		+	+	+	+	+	+	+	
Anemia, macrocytic of pernicious anemia type (nutritional or metabolic) and normocytic metabolic, due to sprue	+	+	+	+	+		+	+	+	+	+	+	+					+	+	+	+	+	+	+	+	+		+	+	
Simple chronic anemias	+	+	+	+			+		+			+	+					+	+	+	+		+	+	+	+		+		
Acquired hemolytic anemia	+	+	+	+	+		+		+		+	+	+	+				+	+	+	+		+	+	+	+	+	+	+	
Sickle cell anemia	+	+	+	+	+		+	+	+		+	+	+	+				+	+	+	+	+	+	+	+	+		+	+	+
Anemia due to acute blood loss	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+		+		
Hypochromic microcytic anemia	+	+	+	+	+		+	+	+	+	+	+	+					+	+	+	+	+	+	+	+	+		+		
Hereditary leptocytosis	+	+	+	+	+		+	+	+			+	+					+	+		+	+	+	+	+	+		+	+	+
Hereditary spherocytosis	+	+	+	+	+		+	+	+		+	+	+	+				+	+	+	+	+	+	+		+		+	+	+
Hemolytic disease of the fetus and newborn	+	+	+			+		+	+		+	+	+				+	+	+						+		+	+	
Hypoprothrombinemia of newborn	+	+	+			+		+	+		+	+	+	+	+			+	+							+		+		
Hemophilia	+		+		+		+	+	+	+	+	+	+	+	+			+	+		+	+	+	+	+	+				
Idiopathic thrombocytopenic purpura	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+	+	+		
Agranulocytosis	+	+	+	+	+		+	+	+		+	+	+					+	+	+	+		+	+	+	+	+	+	+	+
Aplastic (hypoplastic) anemia	+	+	+	+	+		+	+	+		+	+	+					+	+	+	+		+	+	+	+	+	+		
Acute leukemia (myeloblastic, lymphoblastic)	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+	+	+	+	+
Granulocytic leukemia	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+	+	+	+	+
Chronic lymphocytic leukemia	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+	+	+	+	+
Monocytic leukemia	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+		+	+	+
Eosinophilic (basophilic) leukemia	+	+	+	+	+		+	+	+		+	+	+	+	+			+	+	+	+		+	+	+	+	+	+	+	+
Plasma cell myeloma	+	+	+	+			+	+	+		+	+	+					+	+	+	+	+	+	+	+	+		+		
Lipid histiocytosis of kerasin type (Gaucher's disease)	+	+			+		+	+	+		+	+	+					+	+		+	+	+	+	+	+		+		+
Lipid histiocytosis of phosphatide type (Neimann-Pick disease)	+	+			+	+		+	+			+	+			+		+	+			+		+	+	+		+		+
Infectious mononucleosis	+	+	+	+	+		+	+	+		+	+	+	+				+	+	+	+		+	+	+	+		+	+	+
Polycythemia vera	+	+	+	+			+	+	+		+	+	+					+	+	+	+		+	+	+	+		+	+	

* Plus sign indicates that triangular wedge has been punched into space corresponding to given code no. on marginal punched card for that disease.

Figure 13: An example of the information stored on Lipkin and Hardy's marginal punched cards for the diagnosis of hematologic diseases. Martin Lipkin and James Hardy, 'Mechanical Correlation of Data in Differential Diagnosis of Hematological Diseases', *Journal of the American Medical Association* 166, no. 2 (1958): 113-125.

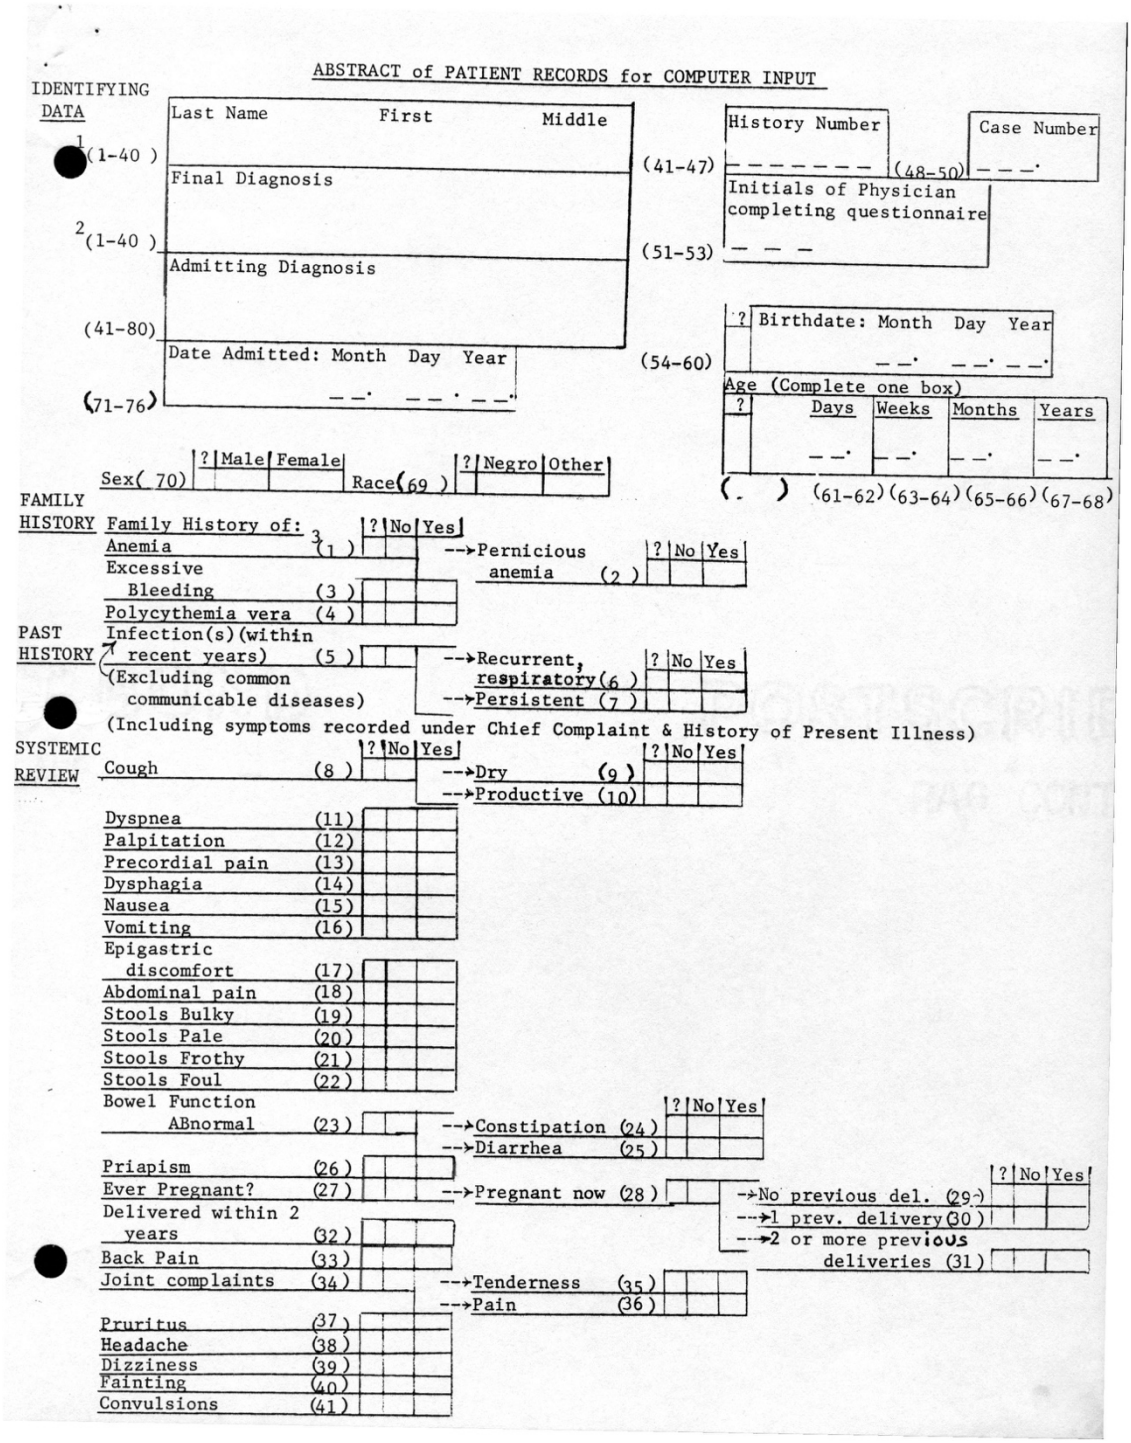

Figure 14: Standardised case history forms were created to impose structure upon medical information for the purposes of computer input. Hagley Museum & Library.

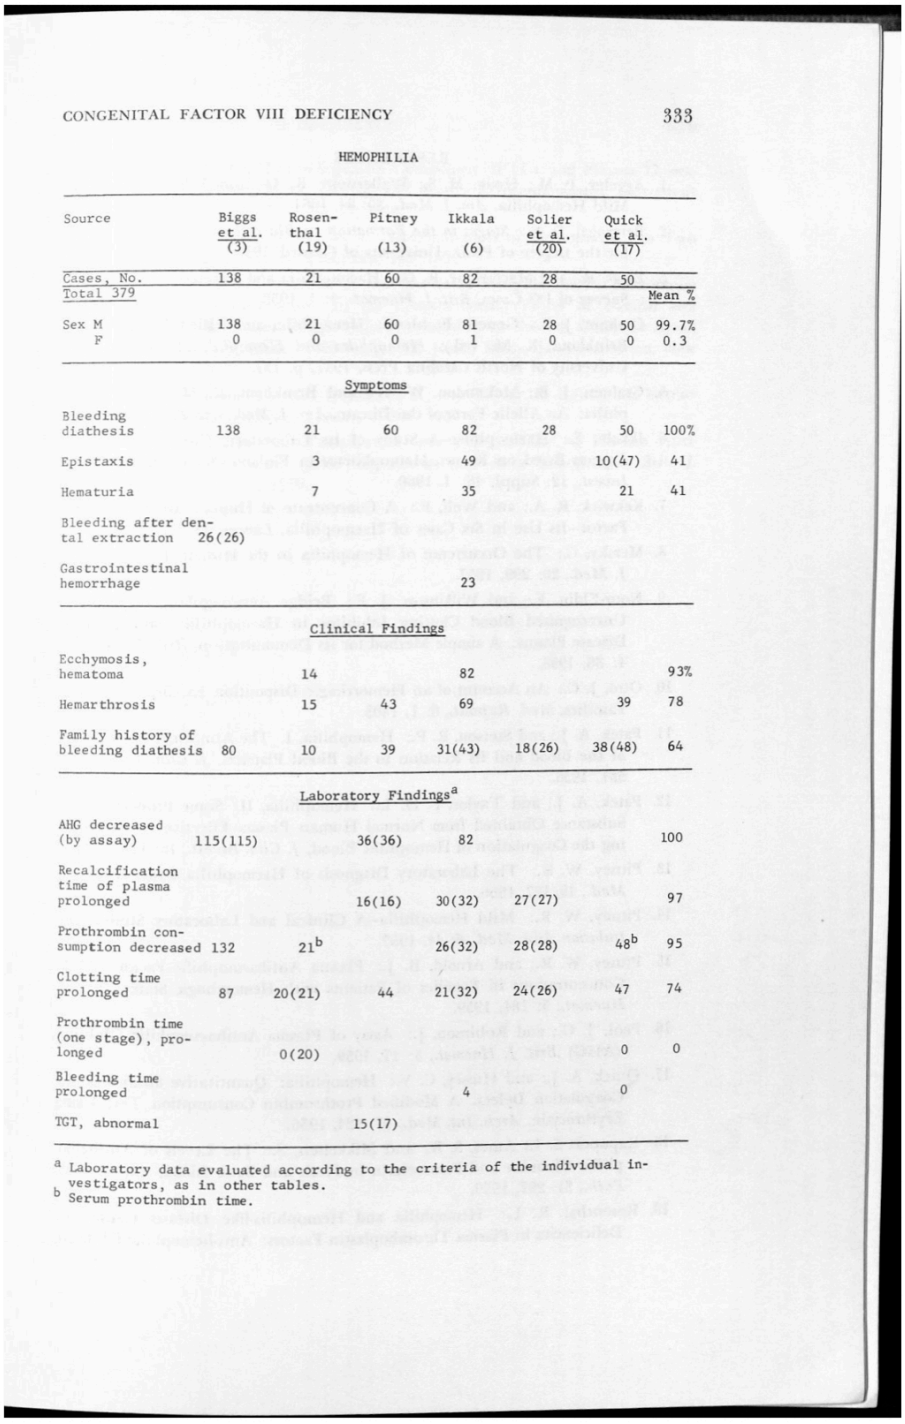

CONGENITAL FACTOR VIII DEFICIENCY 333

HEMOPHILIA

Source	Biggs et al. (3)	Rosen-thal (19)	Pitney (13)	Ikkala (6)	Solier et al. (20)	Quick et al. (17)	
Cases, No. Total 379	138	21	60	82	28	50	Mean %
Sex M	138	21	60	81	28	50	99.7%
F	0	0	0	1	0	0	0.3
Symptoms							
Bleeding diathesis	138	21	60	82	28	50	100%
Epistaxis		3		49		10(47)	41
Hematuria		7		35		21	41
Bleeding after dental extraction	26(26)						
Gastrointestinal hemorrhage				23			
Clinical Findings							
Ecchymosis, hematoma		14		82			93%
Hemarthrosis		15	43	69		39	78
Family history of bleeding diathesis	80	10	39	31(43)	18(26)	38(48)	64
Laboratory Findings[a]							
AHG decreased (by assay)	115(115)		36(36)	82			100
Recalcification time of plasma prolonged			16(16)	30(32)	27(27)		97
Prothrombin consumption decreased	132	21[b]		26(32)	28(28)	48[b]	95
Clotting time prolonged	87	20(21)	44	21(32)	24(26)	47	74
Prothrombin time (one stage), prolonged		0(20)				0	0
Bleeding time prolonged				4		0	
TGT, abnormal			15(17)				

[a] Laboratory data evaluated according to the criteria of the individual investigators, as in other tables.
[b] Serum prothrombin time.

Figure 15: In 1963, Atamer published the findings of his review of the world literature in a volume entitled *Blood Diseases*. The volume included information about different haematologic diseases, including haemophilia (shown above). M. A. Atamer, *Blood Diseases* (New York: Grune & Stratton, 1963).



```
                                FE DEF A

      200000      20       3    MALE
      400000      80       7    FEMALE
      600000     999       2    AGE OF PT. L.T. ONE WEEK
     2200000     999       6    MULTIPARAE
    11000000       6       5    ANORFXIA
    12000000      18       5    BRITTLENESS OF NAILS
    13800500      43       5         DYSPNEA - ON EXERTION
    19200000      50       7    PALLOR (ANEMIA OR RECURRENT ANEMIA)
    21400000       1       5    TONGUE, SORE
    61800000      36       5    CONG HEART DIS/CHR LUNG DIS/A-V ANEURYSM, EVIDENCE OF
    65400000       5       5    HEPATOMEGALY
    69000000      22       7    MASS IN THE ABDOMEN
    70201000       0       5         NEUROLOGICAL ARNORMALITIES - PARAESTHESIA
    71800000      61       7    PERICARDIAL RUB
    77400000      14       5    SPLENOMEGALY
    79000000     999       7    THROMBOSIS, EVIDENCE OF
    79200000      25       5    TONGUE,ATROPHIC OR SMOOTH
   100200000     100       8    A,HGB LTE 12.5 GM. PCT,PCV 35 PCT,RBC 4.10(6) PER CU.MM.
   100200500      98       7    A-HYPOCHROMIC ,MCHC LT 32 PCT/MORPHOLOGICALLY HYPOCHROMIC
   100201500      71       7         ANEMIA - MICROCYTIC (MCV L.T.E. 80 C.U)
   100203000      30       6    A-SEVERE,RBC LTE 2.5X10(6) PER CU.MM.,HGB 7 GM.,PCV 22 PCT
   100800000      70       6    ANISOCYTOSIS, POIKILOCYTOSIS
   101700000     999       6    ERTHROCYTES HYPOCHROMIC
   103400000     999       0    ERYTHSIS,HGB GT 16 GM. PCT,RBC MTE 6X10(6) HEMATCT MT 55
   105400000     999       6    NUCLEATED ERYTHROID CELLS IN PB
   105800000      81       7    PROTOPHORPHYRINS IN ERYTHROCYTES, INCREASED
   106200000     999       6    RETICULOCYTE COUNT,NORMAL
   112600000      11       5    LEUKOPENIA (L.T. 5,000 PER CU.MM.)
   126400000      87       7    BM, HYPERCELLULAR (NORMOBLASTIC)
   127200000     100      10    BM, IRON DECREASED
   155200000      54       6    ACHLORHYDRIA
   158200000      81       7    COPPER (PLASMA),INCREASED
   159200000     100       8    IRON ABSORPTION,INCREASED
   159600000     100      10    IRON BINDING CAPACITY (SERUM),INCREASED
   160200000      93       7    IRON (SERUM),DECREASED
   201400000     100      10    IRON, GOOD RESPONSE

                          NO. OF SYMPTOMS
WEIGHT ONLY                          7
INCIDENCE AND WEIGHT                29
TOTAL----------------               36
```

Figure 16: The definition of iron deficiency anaemia, as stored in the computer. The fourth column represents the collection of items associated with the disease. The third column represents the degree of association between each item and the disease—according to clinical judgement. The second column represents the degree of association between each item and the disease—according to prevalence data. Hagley Museum & Library.

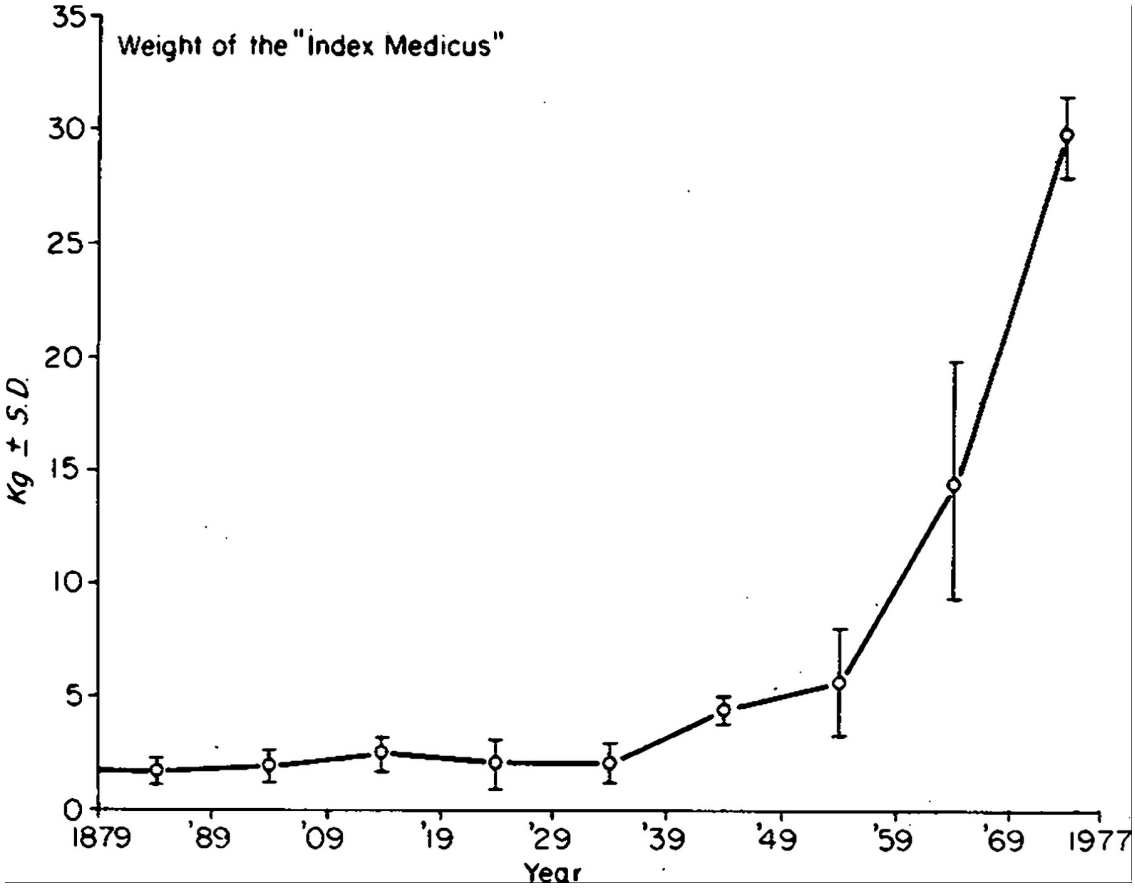

Figure 17: The information explosion in medicine as represented by the weight of the *Index Medicus* from 1879 to 1977. David T. Durack, 'The Weight of Medical Knowledge', *New England Journal of Medicine* 298 (1978): 773-775.



Figure 18: Patient as punched card. Medical Center Archives of New York-Presbyterian/Weill Cornell. 'Computer and Clinical Medicine' *The Sciences* 8 (1968). Medical Center Archives of New York-Presbyterian/Weill Cornell.



Figure 19: Physician as machine. Medical Center Archives of New York-Presbyterian/Weill Cornell.

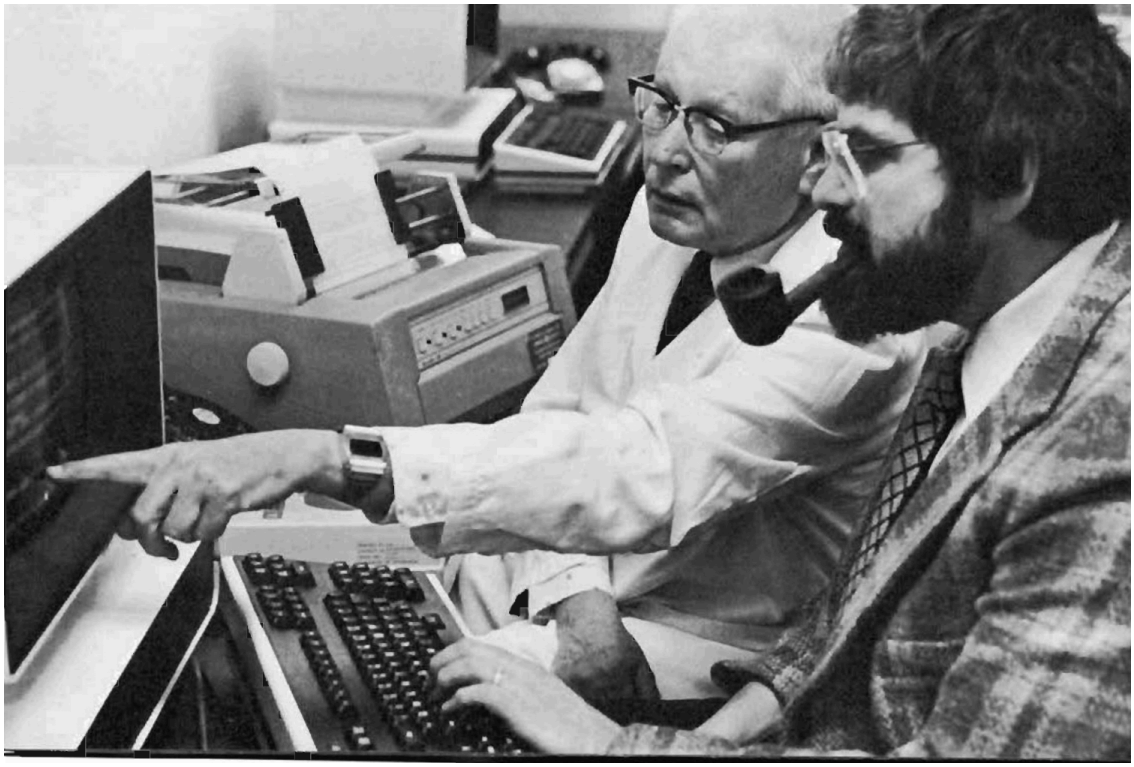

Figure 20: Jack Myers (left) and Harry Pople collaborating together on INTERNIST-I. Gregory Freiherr, *The Seeds of Artificial Intelligence in Medicine: SUMEX-AIM* (Washington, D.C.: U.S. Government Printing Office, 1980), p. 31.



Figure 21: Undated aerial view of 'The Ranch'. Stanford Libraries.

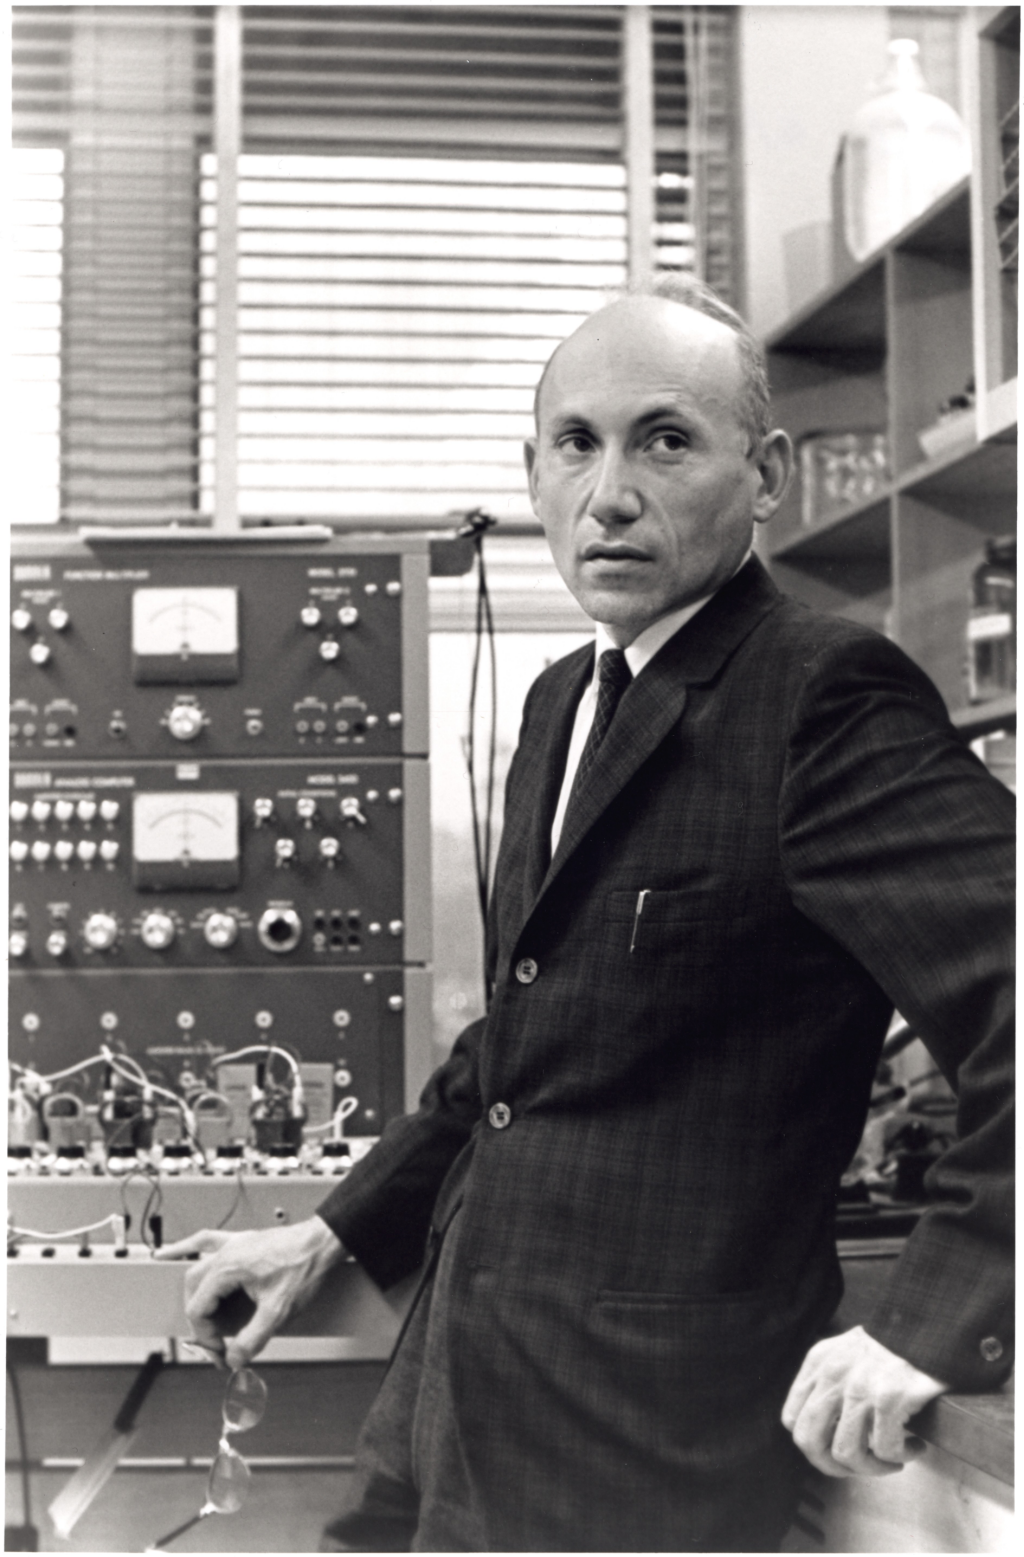

Figure 22: Joshua Lederberg, 1965, in front of exobiology equipment at Stanford. National Library of Medicine.

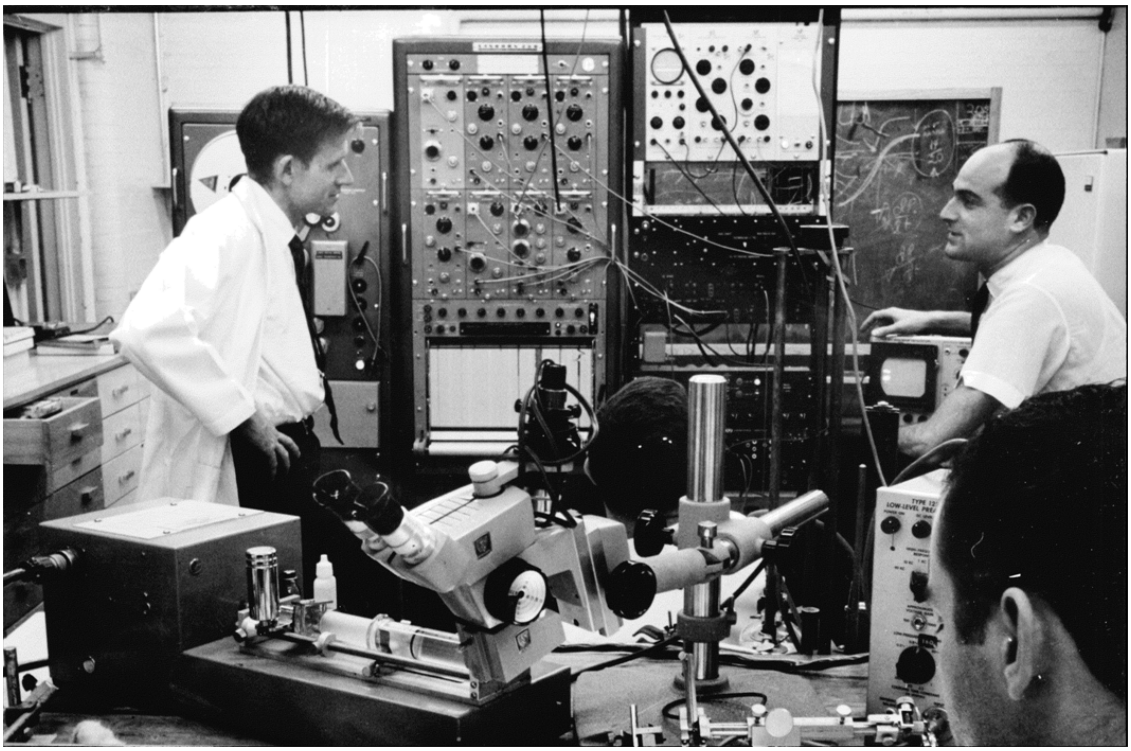

Figure 23: Octo Barnett, circa 1966, standing with colleagues next to a computer at MGH's Laboratory for Computer Science (LCS). Massachusetts General Hospital Laboratory of Computer Science.



Figure 24: A young Edward Shortliffe (right) working with his supervisor and collaborator Bruce Buchanan. Gregory Freiherr, *The Seeds of Artificial Intelligence in Medicine: SUMEX-AIM* (Washington, D.C.: U.S. Government Printing Office, 1980).

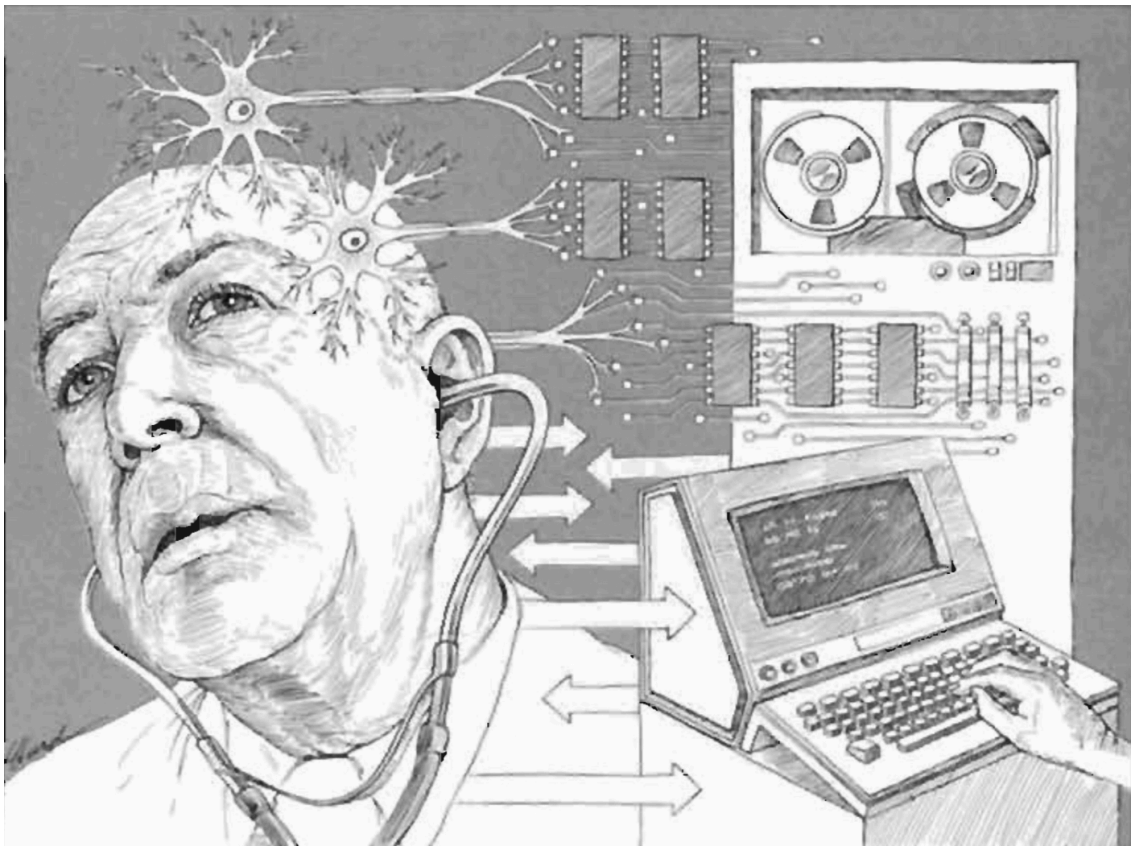

Figure 25: A rendering of how some at the time viewed this process of knowledge extraction. Gregory Freiherr, *The Seeds of Artificial Intelligence in Medicine: SUMEX-AIM* (Washington, D.C.: U.S. Government Printing Office, 1980).

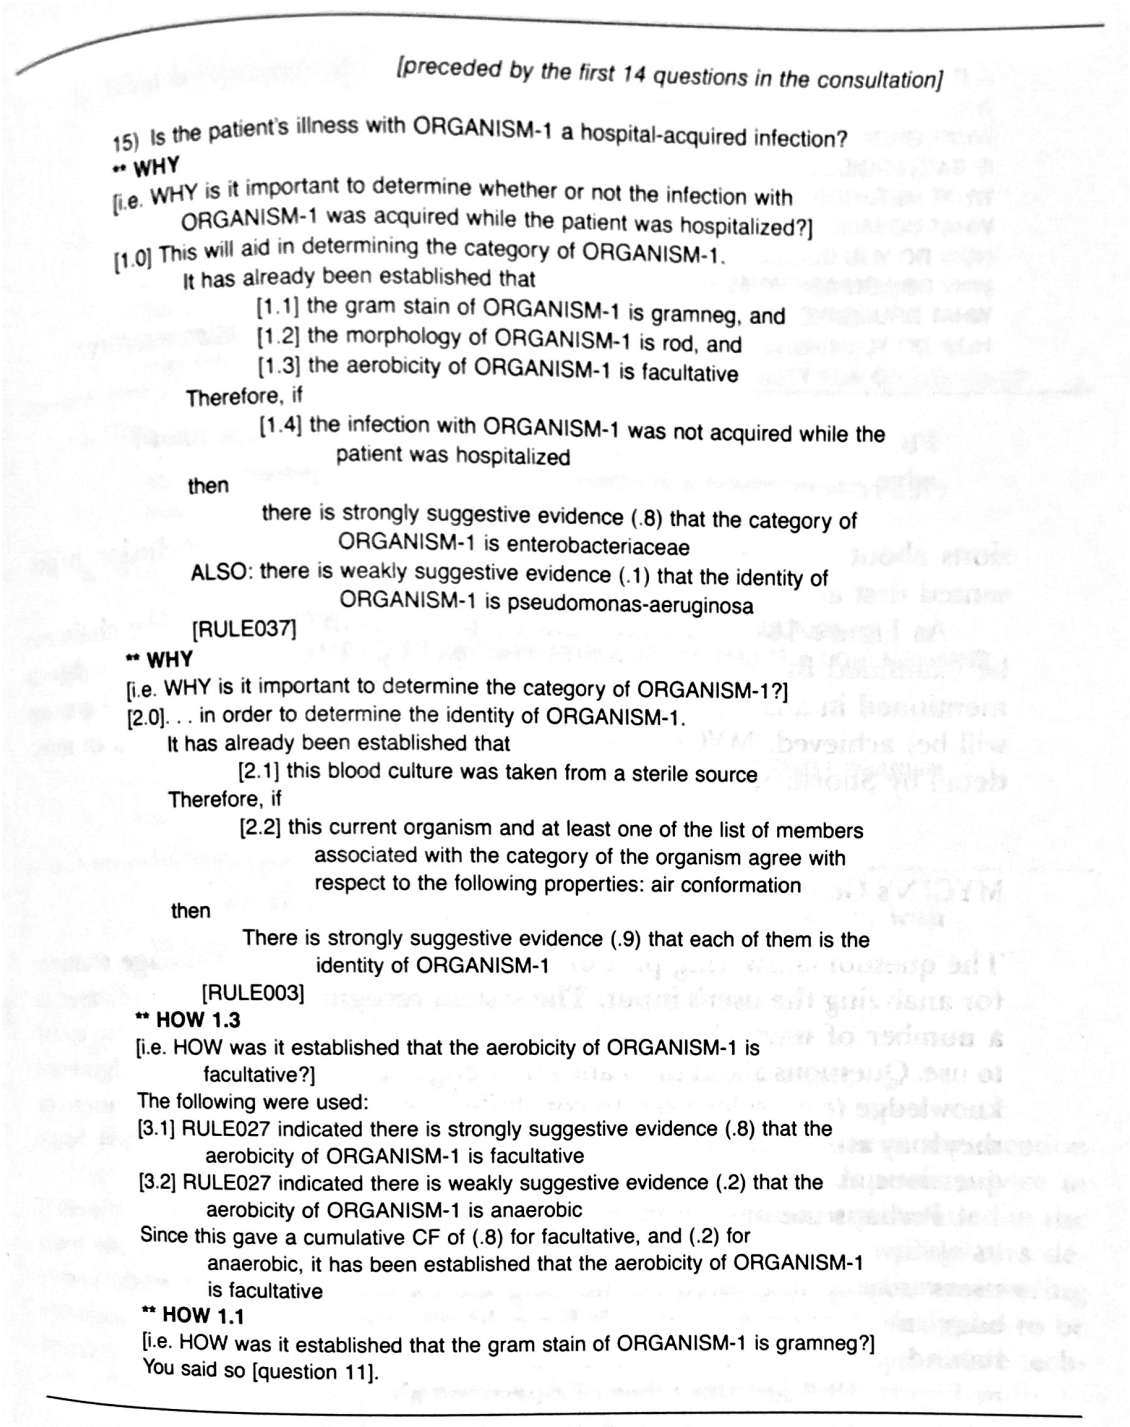

[preceded by the first 14 questions in the consultation]

15) Is the patient's illness with ORGANISM-1 a hospital-acquired infection?
** **WHY**
[i.e. WHY is it important to determine whether or not the infection with ORGANISM-1 was acquired while the patient was hospitalized?]
[1.0] This will aid in determining the category of ORGANISM-1.
It has already been established that
[1.1] the gram stain of ORGANISM-1 is gramneg, and
[1.2] the morphology of ORGANISM-1 is rod, and
[1.3] the aerobicity of ORGANISM-1 is facultative
Therefore, if
[1.4] the infection with ORGANISM-1 was not acquired while the patient was hospitalized
then
there is strongly suggestive evidence (.8) that the category of ORGANISM-1 is enterobacteriaceae
ALSO: there is weakly suggestive evidence (.1) that the identity of ORGANISM-1 is pseudomonas-aeruginosa
[RULE037]
** **WHY**
[i.e. WHY is it important to determine the category of ORGANISM-1?]
[2.0]. . . in order to determine the identity of ORGANISM-1.
It has already been established that
[2.1] this blood culture was taken from a sterile source
Therefore, if
[2.2] this current organism and at least one of the list of members associated with the category of the organism agree with respect to the following properties: air conformation
then
There is strongly suggestive evidence (.9) that each of them is the identity of ORGANISM-1
[RULE003]
** **HOW 1.3**
[i.e. HOW was it established that the aerobicity of ORGANISM-1 is facultative?]
The following were used:
[3.1] RULE027 indicated there is strongly suggestive evidence (.8) that the aerobicity of ORGANISM-1 is facultative
[3.2] RULE027 indicated there is weakly suggestive evidence (.2) that the aerobicity of ORGANISM-1 is anaerobic
Since this gave a cumulative CF of (.8) for facultative, and (.2) for anaerobic, it has been established that the aerobicity of ORGANISM-1 is facultative
** **HOW 1.1**
[i.e. HOW was it established that the gram stain of ORGANISM-1 is gramneg?]
You said so [question 11].

Figure 26: A sample partial consultation with MYCIN. A. Carlisle Scott, William J. Clancey, Randall David, and Edward H. Shortliffe, 'Methods for Generating Explanations', in Bruce G. Buchanan and Edward H. Shortliffe (eds.), *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (Reading, MA: Addison-Wesley, 1984), p. 347.

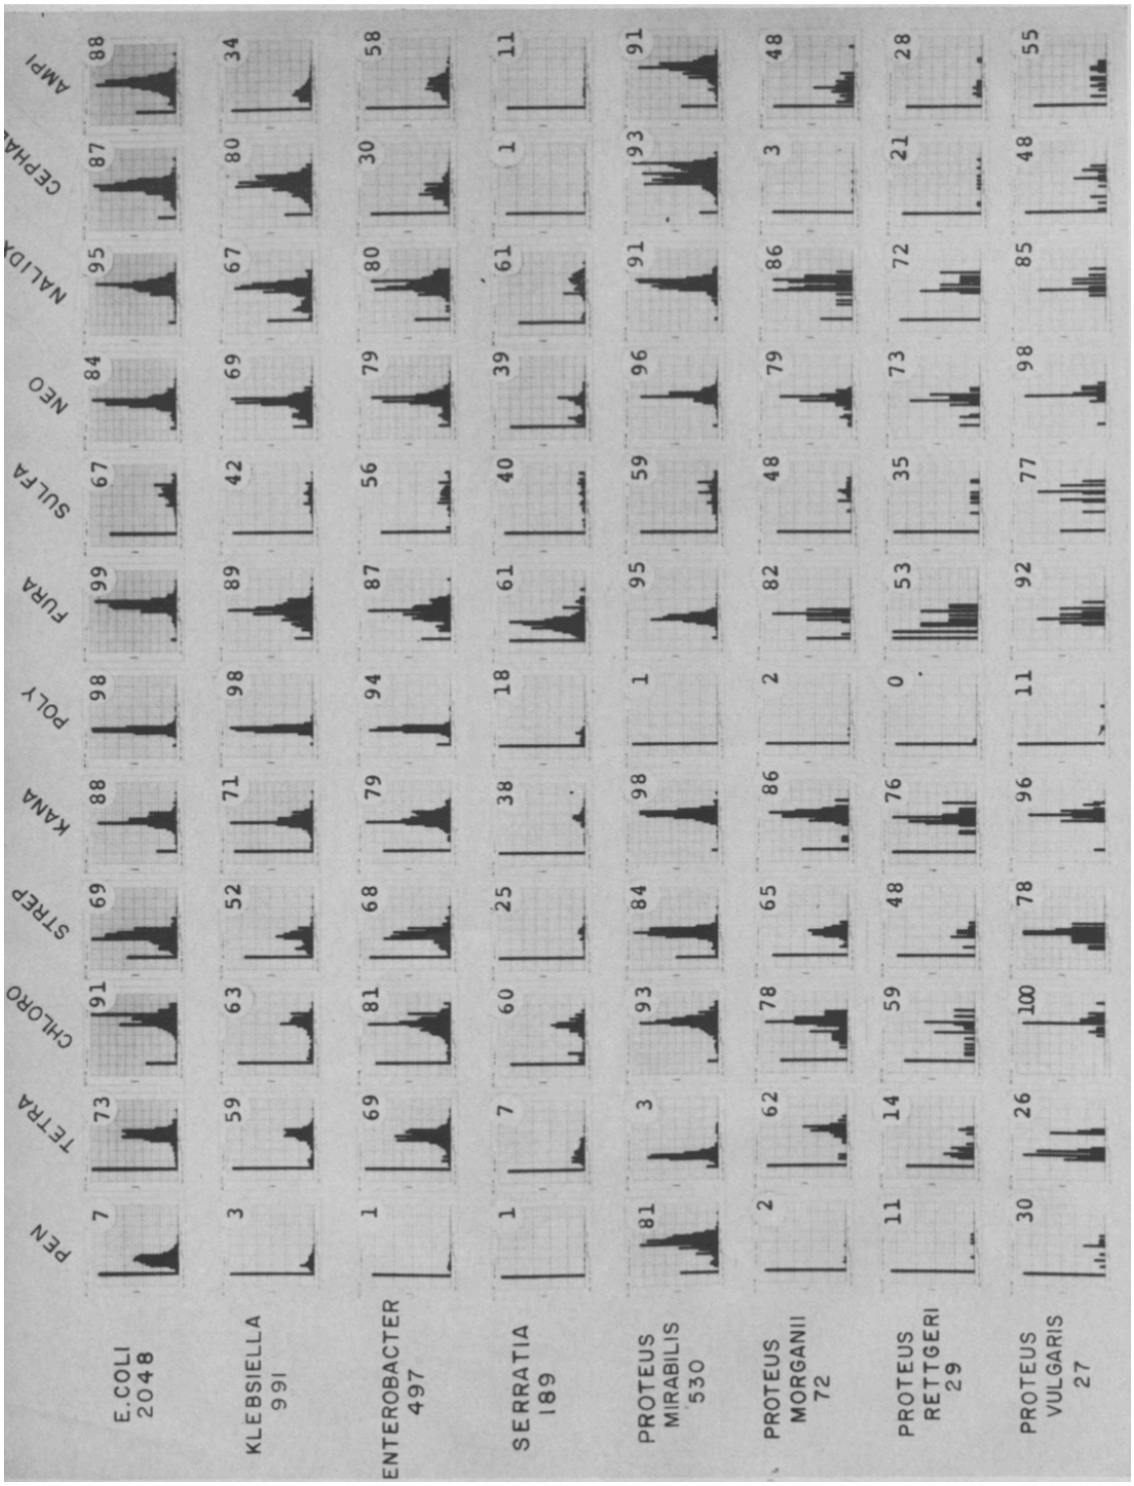

Figure 27: The computer-generated susceptibility plots created by O'Brien and his colleagues. Thomas F. O'Brien, Ralph L. Kent, and Antone A. Medeiros, 'Computer-Generated Plots of Results of Antimicrobial-Susceptibility Tests', *JAMA* 210, no. 1 (1969): 84-92.

BIBLIOGRAPHY

I. Manuscript and Archival Sources

Carnegie Mellon University Libraries (Pittsburgh, PA)
The Herbert Simon Collection

The Hagley Museum and Library (Wilmington, DE)
David Sarnoff Research Center records (Accession 2464.09)
General Research Laboratory records, 1908-1995
Vladimir K. Zworykin papers, 1908-1981

Mayo Clinic Historical Unit, The Mayo Clinic (Rochester, MN)
General and Staff Memoir Collection (MHU-0670)
Mayovox News Collection. 1949-1991 (MHU-0835)
Pamphlet and Booklet Collection (MHU-9992)

Medical Center Archives of New York-Presbyterian/Weill Cornell (New York, NY)
The Keeve Brodman, MD (1906-1979) Papers
The Ralph Engle, MD (1920-2000) Papers
The Harold Wolff, MD (1898-1962) Papers

National Library of Medicine, National Institutes of Health (NIH) (Bethesda, MD)
American Association for Medical Systems and Informatics records (MS C 433)
Octo Barnett papers, 1962-2005 (HMD MS ACC 2006-017)
Marsden Scott Blois Papers 1947-1989 (MS C 583)
Stanley N. Cohen Papers, 1948-2016 (MS C 623)
Morris F. Collen Papers, 1939-2002 (HMD MS ACC 2003-031)
Martin M. Cummings Papers, 1944-1989 (MS C 554)
Lawrence Fagan Papers (ACC 2013-015)
Edward J. Hinman medical informatics collection, 1969-1990 (MS C 547)
Joshua Lederberg Papers, 1904-2008 (MS C 552)
Robert Ledley Papers (HMD MS ACC 2008-036)
Lee Lusted Papers, 1960-1990 (HMD MS ACC 2001-081)
The Regional Medical Programs Collection (MS C 501)
Edward H. Shortliffe Papers (HMD MS ACC 2013-015)
Vernon E. Wilson Papers, 1953-1981 (MS C 492)

The Rockefeller Archive Center (Sleepy Hollow, NY)
Detlev Bronk Papers (FA965)
Rockefeller University records, Catalogs and Directories (FA207)
Rockefeller University Records on the Laboratory of Electronics (FA833)

Special Collections & University Archives, Stanford Libraries (Palo Alto, CA)
Djerassi (Carl) Papers (SC0348)
Feigenbaum (Edward A.) Papers (SC0340)
Stanford University, ACME Project, records (SC0236)

University Archives, University of Minnesota (Minneapolis, MN)

Starke Rosecrans Hathaway papers (uarc 968)

University Archives and Records Management, J. Willard Marriott Library, University of Utah (Salt Lake City, UT)
University of Utah Department of Medical Biophysics and Computing records, 1966-1982 (Acc0196)

II. Newspapers and Periodicals

The American Weekly

The Billboard

The Economist

The Hill

The Hospital Physician

Medical World News

Monterey County Herald

The New Yorker

New York Herald Tribune

New York Times

New York Times Sunday Magazine

The Palm Beach Post

Scientific American

Washington Post

Wired

III. Primary Sources

Abràmoff, Michael D., Philip T. Lavin, Michele Birch, Nilay Shah, and James C. Folk. 'Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices'. *npj Digital Medicine* 1 (2018).

Anderson, J. 'New Self-Administered Medical Questionary'. *British Medical Journal* 4 (1968): 636-638.

Barnett, G. Octo, James J. Cimino, Jon A. Hupp, Edward P. Hoffer. 'Computer Programs to Support Clinical Decision Making – Reply'. *Journal of the American Medical Association* 258, no. 17 (1987): 2376.

Bean, W.B. *Aphorisms from His Bedside Teachings and Writings*. Springfield, IL, 1961.

Brodman, Keeve. 'The Cornell Medical Index-Health Questionnaire: II. As A Diagnostic Instrument'. *Journal of the American Medical Association* 145 (1951): 152–157.

------ 'Diagnostic Decisions by Machine'. *IRE Transactions on Medical Electronics* ME-7 (1960): 216-219.

------ *Men at Work: The Supervisor and His People*. Chicago: Cloud, 1947.

Brodman, Keeve, Albert Erdmann, Jr., Irving Lorge, Jerome Deutschberger, and Harold Wolff. 'The Cornell Medical Index-Health Questionnaire: VII. The Prediction of Psychosomatic and Psychiatric Disabilities in Army Training'. *American Journal of Psychiatry* 111 (1954): 37–40.

Brodman, Keeve, Albert Erdmann, Jr., Irving Lorge, and Charles Gershenson. 'The Cornell Medical Index-Health Questionnaire: III. The Evaluation of Emotional Disturbances'. *Journal of Clinical Psychology* 8 (1952): 119–124

Brodman, Keeve, Albert Erdmann, Jr., Irving Lorge, Charles Gershenson, Harold Wolff, and Todd Broadbent. 'The Cornell Medical Index-Health Questionnaire: IV. The Recognition of Emotional Disturbances in a General Hospital'. *Journal of Clinical Psychology* 8 (1952): 289–293.

Brodman, Keeve, Albert Erdmann, Jr., Irving Lorge, and Harold Wolff. 'The Cornell Medical Index-Health Questionnaire: VI. The Relation of Patients' Complaints to Age, Sex, Race, and Education'. *Journal of Gerontology* 8 (1952): 339–342.

Brodman, Keeve, Albert Erdmann, Jr., Irving Lorge, Harold Wolff, and Todd Broadbent. 'The Cornell Medical Index: An Adjunct to Medical Interview'. *Journal of the American Medical Association* 140 (1949): 530–534.

Brodman, Keeve and Adrianus van Woerkom. 'Computer-Aided Diagnostic Screening for 100 Common Disease'. *Journal of the American Medical Association* 197 (1966): 179-183.

Brodman, Keeve, Adrians van Woerkom, Albert Erdmann, Jr., and Leo Goldstein. 'Interpretation of Symptoms with a Data-Processing Machine'. *A.M.A. Archives of Internal Medicine* 103 (1959): 116-122.

Buchanan, Bruce and Edward Feigenbaum. 'Dendral and Meta-Dendral: Their Applications Dimension'. *Artificial Intelligence* 11 (1978): 5-24.

------ ‘Dendral and Meta-Dendral: Roots of Konwledge Systems and Expert System Applications’. *Artificial Intelligence* 59 (1993): 233-240.

Cabot, Richard C. and Hugh Cabot. ‘Ante-mortem and Post-mortem Records as Used in the Weekly Clinico-pathological Exercises. Case 10271’. *Boston Medical and Surgical Journal* 191, no. 4 (1924): 30-34.

Carbonell, Jamie R. ‘AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction’. *IEEE Transactions on Man-Machine Systems* 11, no. 4 (1970): 190-202.

Castleman, Benjamin and H. Robert Dudley. *Clinicopathological Conference of the Massachusetts General Hospital, Selected Medical Cases*. Boston: Little, Brown and Co., 1960.

‘Challenging the Leader’. *Time* 85, no. 5 (1965): 71.

Clancey, William J. ‘Epistemology of a Rule-Based Expert System’. *Artificial Intelligence* 20, no. 3 (1983): 215-251.

------ ‘Notes on “Epistemology of a Rule-Based Expert System”’. *Artificial Intelligence* 59, no. 1-2 (1993): 197-204.

Clouser, K. Danner. ‘Approaching the Logic of Diagnosis’. In Kenneth F. Schaffner. *Logic of Discovery and Diagnosis in Medicine*. Berkeley, CA: University of California Press, 1985, 35-55.

Cochrane, A. L. P. J. Chapman, and P. D. Oldham. ‘“Observers” Errors in Taking Medical Histories’. *Lancet* (5 May 1951): 1007-1009.

Cohen, Stanley N. *Drug Interactions: A Handbook for Clinical Use*. Baltimore: Williams and Wilkins, 1974.

Cohen, Stanley N., Marsha F. Armstrong, Linda Crouse, and Gilbert S. Hunn. ‘A Computer-Based System for Prospective Detection and Prevention of Drug Interactions’. *Drug Information Journal* 72 (1972): 81-86.

Crosby, William H. ‘Chess and Combat: The Algorithm in Medicine’. *Journal of the American Medical Association* 238, no. 25 (1977): 2721.

Cully, Megan. ‘The Politics of Antibiotic Use’. *Nature* 509 (2014): S16-S17.

‘Cytoanalyzer’. *Public Health Reports* 72, no. 11 (1957): 1038.

de Sollar, Derek J. *Little Science, Big Science*. New York and London: Columbia University Press, 1963.

Duncomb, Raynor and Morris Davis. ‘Obituary: A. J. J. van Woerkom, 1915-1991’. *Bulletin of the American Astronomical Society* 23 (1991): 1495.

Durack, David. ‘The Weight of Medical Knowledge’. *New England Journal of Medicine* 298, no. 14 (1978): 773-775.

Ebald R. and R. Lane. ‘Digital Computers and Medical Logic’. *IRE Transactions of Medical Electronics* 7 (1960): 283-288.

Eddy, David and Charles Clanton. ‘The Art of Diagnosis: Solving the Clinicopathological Exercise’. *New England Journal of Medicine* 306, no. 21 (1982): 1263-1268.

Elfiky, Aymen A., Maximilian J. Pany, Ravi B. Parikh, and Ziad Obermeyer. ‘Development and Application of a Machine Learning Approach to Assess Short-term Mortality Risk Among Patients with Cancer Starting Chemotherapy’. *JAMA Network Open* 1, no. 3 (2018): e180926.

Engelmore, Robert S. ‘Artificial Intelligence and Knowledge Based Systems: Origins, Methods and Opportunities for NDE’. In D. O. Thompson and D. E. Chimenti. *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 6. Boston: Springer, 1987, 1-20.

Engle, Ralph L. ‘Attempts to Use Computers as Diagnostic Aids in Medical Decision Making: A Thirty Year Experience’. *Perspectives in Biology and Medicine* 35, no. 2 (1992): 207-219.

------ ‘Medical Diagnosis: Present, Past, and Future. II. Philosophical Foundations and Historical Development of Our Concepts of Health, Disease, and Diagnosis’. *Archives of Internal Medicine* 112 (1963): 520-529.

------ ‘Medical Diagnosis: Present, Past, and Future. III. Diagnosis in the Future, Including a Critique on the Use of Electronic Computers as Diagnostic Aids to the Physician’. *Archives of Internal Medicine* 112 (1963): 530-543.

Engle, Ralph L. and B. J. Davis. ‘Medical Diagnosis: Present, Past, and Future I. Present Concepts of the Meaning and Limitations of Medical Diagnosis’. *Archives of Internal Medicine* 112, no. 4 (1963): 512-519.

Engle, Ralph L. and Betty J. Flehinger. ‘HEME2: A Long History of Computer-Aided Hematologic Diagnosis’. In Bruce I. Blum. *Proceedings of the Sixth Annual Symposium on Computer Applications in Medical Care*. Computer Society Press, 1982, 763-767.

Engle, Ralph L., Betty J. Flehinger, Scott Allen, Richard Friedman, Martin Lipkin, B. J. Davis, and Leo Leveridge. ‘HEME: A Computer Aid to Diagnosis of Hematologic Disease’. *Bulletin of the New York Academy of Sciences* 52, no. 5 (1976): 584-600.

Erdmann, Jr. Albert,, Keeve Brodman, Irving Lorge, and Harold Wolff. ‘The Cornell Medical Index-Health Questionnaire: V. Outpatient Admitting Department of a General Hospital’. *Journal of the American Medical Association* 149 (1952): 550–551.

Estes, W. K. ‘Of Models and Men’. *American Psychologist* 12, no. 10 (1957): 609-617.

Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 'Dermatologist-level classification of skin cancer with deep neural networks'. *Nature* 542 (2017): 115–118.

Feinstein, Alvan. 'Clinical Epidemiology I: The Population Experiments of Nature and of Man in Human Illness'. *Annals of Internal Medicine* 69 (1968): 807-820.

------ *Clinical Judgement*. Baltimore: Williams and Wilkins, 1967.

Flehinger, Betty J. and Ralph L. Engle. 'HEME: A Self-Improving Computer Program for Diagnosis-Oriented Analysis of Hematologic Diseases'. *IBM Journal of Research and Development* 19, no. 6 (1975): 557-564.

Fowler, Raymond D. 'Computer Interpretation of Personality Tests: The Automated Psychologist'. *Comprehensive Psychology* 8, no. 6 (1967): 455-467.

Freiherr, Gregory. 'MYCIN: A Medical and Educational Aid'. *Research Resources Reporter* 3, no. 12 (1979): 1-6.

------ *The Seeds of Artificial Intelligence in Medicine: SUMEX-AIM*. Washington, D.C.: U.S. Government Printing Office, 1980.

Gao, Daozhou, Thomas M. Lietman, and Travis C. Porco. 'Antibiotic Resistance as Collaterol Damage: The Tragedy of the Commons in a Two-Disease Setting'. *Mathematical Biosciences* 263 (2015): 121-132.

Geelhoed, G. W. and E. M. Druy. 'Management of the Adrenal "Incidentaloma"'. *Surgery* 92 (1982): 866-874.

Getting, Vlado A. and Herbert L. Lombard, 'The Evaluation of Pilot Clinics—The Mass Screening or Health-Protection Program'. *New England Journal of Medicine* 247 (1952): 460-465.

Glaser, E. M. and G. C. Whittow, 'Experimental Errors in Clinical Trials'. *Clinical Science* 13, no. 2 (1954): 199-210.

Gledhill, V. X., J. D. Mathews, and I. R. Mackay. 'Computer-Aided Diagnosis: A Study of Bronchitis'. *Methods of Information in Medicine* 11 (1972): 228-233.

Gorry, G. Anthony. 'Computer-Assisted Clinical Decision Making'. *Methods of Information in Medicine* 12 (1973): 45-51.

Gorry, G. Anthony and G. Octo Barnett. 'Experience with a Model of Sequential Diagnosis'. *Computers and Biomedical Research* 1 (1968): 490-507.

Grant, Graham, and R.A.N. Hitchens. 'Correspondence: The Patient Tells His Story'. *British Medical Journal* 4801 (10. Jan. 1953): 100-101.

Greenberg, Selig. ‘The Decline of the Healing Art’. *Harper’s Magazine* 221 (1960): 132-150.

Greenes, Robert A., Bruce G. Buchanan, and Donald Ellison. ‘Presentation of the 2006 Morris F. Collen Award to Edward H. (Ted) Shortliffe’. *Journal of the American Medical Informatics Association* 14, no. 3 (2007): 376-385.

Harris, Daniel K. ‘Computer-Assisted Decision Support: A New Available Reality’. *Journal of the American Medical Association* 258, no. 1 (1987): 86.

------ ‘Computer Programs to Support Clinical Decision Making – Reply’. *Journal of the American Medical Association* 258, no. 17 (1987): 2376.

Harris, Jesse G., Jr. ‘Judgmental versus Mathematical Prediction: An Investigation by Analogy of the Clinical versus Statistical Controversy’. *Behavioral Science* 8, no. 4 (1963): 324-335.

Heckerman, David and Eric Horvitz. ‘The Myth of Modularity in Rule-Based Systems for Reasoning with Uncertainty’. *Machine Intelligence and Pattern Recognition* 2 (1986): 23-34.

Hickam, David H., Edward H. Shortliffe, Miriam B. Bischoff, A. Carlisle Scott, and Charlotte D. Jacobs. ‘The Treatment Advice of a Computer-Based Cancer Chemotherapy Protocol Advisor’. *Annals of Internal Medicine* 103, no. 6 (1985): 928-936.

Hodgins, Eric. ‘Listen: The Patient’. *New England Journal of Medicine* 274 (1966): 657-661.

Hoefler, Don D. ‘Silicon Valley in the USA’. *Electronic News*, 11 Jan. 1971.

Hollis, Aidan and Peter Maybarduk. ‘Antibiotic Resistance is a Tragedy of the Commons That Necessitates Global Cooperation’. *Journal of Law, Medicine and Ethics* 43 (2015): 33-37.

Holm, Elizabeth. ‘In Defense of the Black Box’. *Science* 364, no. 6435 (2019): 26-27.

Horvitz, Erc and David Heckerman. ‘Modular Belief Updates and Confusion about Measures of Certainty in Artificial Intelligence Research’. In Laveen Kanal and John Lemmer. *Proceedings of the First Conference on Uncertainty in Artificial Intelligence*. Arlington, VA: AUAI Press, 1985, 283-286.

Hubbard, William N. ‘Welcome and Prologue’. In John A. Jacquez. *The Diagnostic Process: Proceedings of a Conference Held at the University of Michigan, May 9-11, 1963*. April 1964, 3-5.

Ingelfinger, F. J ‘Algorithms, Anyone?’. *New England Journal of Medicine* 288 (1973): 847-848.

Jacquez, John A. ‘Preface’. In John A. Jacquez. *The Diagnostic Process: Proceedings of a Conference Held at the University of Michigan, May 9-11, 1963*. April 1964, iii-iv.

Jones, Richard S. ‘Letters: Use of Computers in Medical Practice’. *Canadian Medical Association Journal* 122, no. 4 (1980): 400.

Karlgren, Hans. Review of *Computational Linguistic in Medicine*, Werner Schneider and Anna-Lena Sågvall Hein. *Journal of Pragmatics* 4, no. 2 (1980): 183.

Kent, Daniel L., Edward H. Shortliffe, Robert W. Carlson, Miriam B. Bischoff, Charlotte D. Jacobs. ‘Improvements in Data Collection Through Physician Use of a Computer-Based Chemotherapy Treatment Consultant’. *Journal of Clinical Oncology* 3 (1985): 1409-1417.

Kleinmuntz, Benjamin. ‘The Processing of Clinical Information by Man and Machine’. In Benjamin Kleinmuntz. *Formal Representation of Human Judgement*. New York: John Wiley, 1968.

Knill-Jones, R. P., R. B. Stern, D. H. Grimes, J. D. Maxwell. R. P. H. Thompson, and R. Williams. ‘Use of a Sequential Bayesian Model in Diagnosis of Jaundice by Computer’. *British Medical Journal* 1 (1973): 530-533.

Kulikowski, Casimir. ‘Knowledge-Based Systems in Biomedicine: A 10 Year Retrospective’. *Proceedings of the Annual Symposium on Computer Application in Medical Care* (1986): 423-424.

Kunin, C. M., T. Tupasi, and W. A. Craig. ‘Use of Antibiotics: A Brief Exposition of the Problem and Some Tentative Solutions’. *Annals of Internal Medicine* 79 (1973): 555-560.

Lakhani, P. and B. Sundaram. ‘Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks’. *Radiology* 284 (2017): 574–582.

Leavell, Byrd and Oscar Thorup, Jr. *Fundamentals of Clinical Hematology*. Philadelphia and London: W. B. Saunders Company, 1960.

Ledley, Robert S. and Lee B. Lusted. ‘Reasoning Foundations of Medical Diagnosis’. *Science* 130, no. 3366 (1959): 9-21.

Lewis, Charles E. ‘Variations in the Incidence of Surgery’. *New England Journal of Medicine* 281 (1969): 880-884.

Levin, James E. ‘Computer Programs to Support Clinical Decision Making’. *Journal of the American Medical Association* 258, no. 17 (1987): 2375.

Lipkin, Martin. ‘Correlation of Data with a Digital Computer in the Differential Diagnosis of Heamtologic Disease’. *IRE Transactions on Medical Electronics* ME-7, no. 4 (1960): 243-246.

------ ‘Data Processing in the Diagnostic Process’. In John A. Jacquez. *The Diagnostic Process: Proceedings of a Conference Held at the University of Michigan, May 9-11, 1963*. Apr. 1964, 255-280.

------ ‘Digital and Analogue Computer Methods Combined to Aid in the Differential Diagnosis of Heamtological Diseases’. *Circulation Research* 11 (1962): 607-613.

Lipkin, Martin, Ralph L. Engle, Jr., B. J. Davis, Vladimir Zworykin, R. Ebald, M. Sendrow, and Carl Berkley. ‘Digital Computer as Aid to Differential Diagnosis: Use in Hematologic Diseases’. *Archives of Internal Medicine* 108, no. 1 (1961): 56-72.

Lipkin, Martin, Ralph L. Engle, Jr., Betty Flehinger, Louis Gerstman, and M. A. Atamer. ‘Computer-Aided Differential Diagnosis of Hematologic Diseases’. *Annals of the New York Academy of Sciences* 161 (1969): 670-679.

Lipkin, Martin and James Hardy. ‘Differential Diagnosis of Hematologic Diseases Aided by Mechanical Correlation of Data’. *Science* 125, no. 3247 (1957): 551-552.

------ ‘Mechanical Correlation of Data in Differential Diagnosis of Hematologic Diseases’. *Journal of the American Medical Asssocaition* 166, no. 2 (1958): 113-125.

Lipkin, Martin and Max Woodbury. ‘Analytical Studies Related to the Differential Diagnosis of Hematologic Disease’. *Blut* 9, no. 7 (1963): 449-454.

------ ‘Coding of Medical Case History Data for Computer Analysis’. *Communications of the ACM* 5, no. 10 (1962): 532-534.

Litjens, G., C. Sánchez, N. Timofeeva, M. Hermsen, I. Nagtegaal, I. Kovacs, C. Hulsbergen-van de Kaa, P. Bult, B. van Ginneken, and J. van der Laak. ‘Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis’. *Science Reports* 6 (2016): 26286.

Lodwick, G. S., C. L. Haun, W. E. Smith, R. F. Keller, and E. D. Robertson. ‘Computer Diagnosis of Primary Bone Tumors: A Preliminary Report’. *Radiology* 80 (1963).

London, Alex John. ‘Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability’. *The Hastings Center Report* 49, no. 1 (2019): 15-21.

Lykken, D.T. ‘A Method of Actuarial Pattern Analysis’. *Psychological Bulletin* 53 (1956): 102-107.

Mann, Richard D. ‘A Critique of P. E. Meehl’s *Clinical versus Statistical Prediction*’. *Behavioral Science* 1, no. 3 (1956): 224-230.

McCarthy, John, Marvin Minkey, Nathaniel Rochester, Claude Shannon. ‘A Proposal For the Dartmouth Summer Research Project on Artificial Intelligence’. *AI Magazine* 27 (2006): 12-14.

McCulloch, Warren S. and Walter Pitts. ‘How We Know Universals: The Perception of Auditory and Visual Forms’. *Bulletin of Mathematical Biophysics* 9 (1947): 127-147.

McTernan, Edmund and Dean Crocker. ‘Push-button medicine is no pipe dream!’. *Hospital Physician* (1969): 85.

Meehl, Paul E. *Clinical versus Statistical Prediction: A Theoretical Analysis and Review of the Evidence*. Minneapolis: University of Minnesota Press, 1954.

------ ‘A Comparison of Clinicians with Five Statistical Methods of Identifying Psychotic MMPI Profiles’. *Journal of Counseling Psychology* 6, no. 2 (1959): 102-109.

------ ‘Wanted—A Good Cookbook’. *American Psychologist* 11, no. 6 (1956): 263-272.

Meehl, Paul E. and W.G. Dahlstrom. ‘Objective Configural Rules for Discriminating Psychoatic from Neurotic MMPI Profiles’. *Journal of Consulting Psychology* 24 (1960): 375-387.

Mittelmann, Bela and Keeve Brodman, ‘The Cornell Indices and the Cornell Word Form: 1. Construction and Standardization’. *Annals of the New York Academy of Sciences* 46 (1946): 573–578.

Moser, Robert. *Diseases of Medical Progress*, 3rd edition. Springfield: C.C. Thomas, 1969.

Myers, Jack D. ‘The Background of INTERNIST I and QMR’. In *HMI ’87 Proceedings of ACM Conference on History of Medical Informatics*. New York: ACM, 1987, 195-197.

Nash, F. A. ‘Diagnostic Reasoning and the Logoscope’. *Lancet* 276, no. 7166 (1960):1442-1446.

National Commission on Technology, Automation and Economics Progress. *Technology and the American Economy*, vol. 1. Washington, D.C.: U.S. Government Printing Office, 1966.

Newell, Allen. ‘Forward’. In Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley, 1984, xi-xvi.

Newell, Allen, C. J. Shaw, and Herbert Simon. ‘Elements of a Theory of Human Problem Solving’. *Psychological Review* 65, no. 3 (1958): 151-166.

Obermeyer, Ziad, and Ezekiel J. Emanuel. ‘Predicting the Future—Big Data, Machine Learning, and Clinical Medicine’. *New England Journal of Medicine* 375 (2016): 1216-1219.

Obermeyer, Ziad, and Thomas H. Lee. ‘Lost in Thought—The Limits of the Human Mind and the Future of Medicine’. *New England Journal of Medicine* 377 (2017): 1209-1211.

O'Brien, Thomas F., Ralph L. Kent, and Antone A. Medeiros. 'Computer -Generated Plots of Rusults of Antimicrobial-Susceptibility Tests'. *Journal of the American Medical Association* 210, no. 1 (1969); 84-92.

Paxon, Harry T. 'The Computer: A Report on How It's Affecting the Hospital Physicians'. *Hospital Physician* (Sep. 1966): 37.

'The Patient Tells His Story'. *British Medical Journal* 4796 (1952): 1246-1247.

Physicians for a Growing American. Report of the Surgeon General's Group of Medical Education, publication No. 709. Besthesda, MD: U.S. Department of Health, Education, and Welfare, 1959.

Pearson, R. John, Bjorn Smedby, Rognard Berfenstam, Robert Logan, Alex Burgess, Jr., and Osler Peterson. 'Hospital Caseloads in Liverpool, New England, and Uppsala: An International Comparison'. *Lancet* 292, no. 7567 (1968): 559-566.

Philips, O.C., T. M. Frazier, T. D. Graff, and T. J. DeKornfeld. 'The Baltimore Anesthesia Study Committee: Review of 1,024 Postoperative Deaths'. *Journal of the American Medical Association* 174 (1960): 2015-2084.

Pickering, George. 'Medicine and Society—Past, Present, and Future'. *British Medical Journal* 1, no. 23 (1971): 191-196.

Pickworth, F. A. 'Correspondence: The Patient Tells History Story'. *British Medical Journal* (27 Dec. 1952): 1414.

The President's Commission on the Health Needs of the Nation. *Building America's Health: A Report to the President*. Washington, D.C.: U.S. Government Printing Office, 1952.

The President's Commission on Heart Disease, Cancer and Stroke. *Report to the President: A Program to Conquer Heart Disease, Cancer and Stroke* vol. 1. Washington, D.C.: U.S. Government Printing Office, 1964.

'Questionaires in Clinical Trials'. *British Medical Journal* 4875 (12 Jun. 1954): 1366-1367.

Redelmeier, Donald A. and Amos Tversky. 'Discrepancy between Medical Decisions for Individual Patients and for Groups'. *New England Journal of Medicine* 322 (1990): 1162-1164.

Reik, Theodor. *Listening with the Third Ear*. New York: Farrar & Straus, 1948.

'Report on the Screening Clinic'. *New England Journal of Medicine* 243 (1950): 275.

Reztak, K. E. and R. B. Williams. 'A Review of Antibiotic Therapy in Patients with Systematic Infections'. *American Journal of Hopsital Pharmacy* 29 (1972): 935-941.

Rikli, Arthur E., Fred V. Lucas, and Fred Frazier. 'The Automated Physician's Assistant'. *Clinical Engineering News* 3, no. 4 (1975): 3-5.

Roberts, A. W. and J. A. Visconti. 'The Rational and Irrational use of Systematic Antimicrobial Drugs'. *American Journal of Hospital Pharmacy* 29 (1972): 828-834.

Robins, Eli. 'Categories Versus Dimensions in Psychiatric Classication'. *Psychiatic Annals* 6, no. 8 (1976): 39, 42-43, 46-47, 51, 55.

Sarbin, Theodore. 'Clinical Psychology—Art or Science?'. *Psychometrika* 6 (1941): 391-399.

------ 'A Contribution to the Study of Actuarial and Individual Methods of Prediction'. *American Journal of Sociology* 48 (1943): 593-602.

------ 'The Logic of Prediction in Psychology'. *Psychological Review* 51 (1944): 210-228.

------ 'The Relative Accurarcy of Clinical and Statistical Predictions of Academic Achievement'. *Psychological Bulletin* 38 (1941): 714.

Savage, Leonard J. *The Foundations of Statistics*. New York, NY: John Wiley & Sons, Inc., 1954.

Schapira, Max, Edith Kepes, and Elliott Hurwitt. 'An Analysis of Deaths in the Operating Room and within 24 Hours of Sugery'. *Anesthesia and Analgesia* 39 (1960): 149-157.

Scheckler, W. E. and J. V. Bennett. 'Antibiotic Usage in Seven Community Hospitals'. *Journal of the American Medical Association* 213 (1970): 264-267.

Schimmel, Elihu M. 'The Hazards of Hospitalization'. *Annals of Internal Medicine* 60, no. 1 (1964): 100-110.

Scotch, Norman A. and H. Jack Geiger. 'An Index of Symptom and Disease in Zulu Culture'. *Human Organization* 22, no. 4 (1963-1964): 304-311.

Selfridge, Oliver. 'Pandemonium: A Paradigm of Learning'. In National Laboratory. *Mechanisation of Thought Processes*. London: Her Majesty's Stationary Office, 1959, 511-529.

Sheiner, Lewis B., George Brecher, and Lawrence A. Wheeler. 'Clinical Laboratory Use in the Evaluation of Anemia'. *Journal of the American Mediccal Association* 238, no. 25 (1977): 2709-2714.

------ 'Letters: The Algorithm in Medicine'. *Journal of the American Medical Association* 239, no. 19 (1978): 1957.

Shortliffe, Edward H. 'The Changing Nature of Telecommunications and Information Infrastructure for Health Care'. In *The Changing Nature of*

Telecommunications/Information Infrastructure. Washington, D.C.: National Academy Press, 1995, 67-73.

------. 'Clinical Consultation Systems: Designing for the Physician As Computer User'. *Proceeding of the Annual Symposium on Computer Applications in Medical Care* (1981): 235-236.

------ *Computer-Based Medical Consultations: MYCIN*. New York: Elsevier, 1976.

------ 'Computer Programs to Support Clinical Decision Making'. *Journal of the American Medical Association* 258, no. 1 (1987): 61-66.

------ 'Computer Programs to Support Clinical Decision Making – Reply'. *Journal of the American Medical Association* 258, no. 17 (1987): 2375-2376.

------ 'Medical Expert Systems-Knowledge Tools for Physicians'. *Western Journal of Medicine* 145, no. 6 (1986): 830-839.

------ 'MYCIN: A Knowledge-Based Computer Program Applied to Infectious Diseases'. *Proceeding of the Annual Symposium on Computer Applications in Medical Care* (1977): 66-69.

Shortliffe, Edward H. and Bruce G. Buchanan. 'A Model of Inexact Reasoning in Medicine'. *Mathematical Biosciences* 23 (1975): 351-379.

Shortliffe, Edward H., Randall Davis, Stanton G. Axline, Bruce G. Buchanan, C. Cordell Green, and Stanley N. Cohen. 'Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system'. *Computers in Biomedical Research* 8, no. 4 (1975): 303-320.

Shortliffe, Edward H., A. Carlisle Scott, Miriam B. Bischoff, A. Bruce Campbell, William van Melle, and Charlotte D. Jacobs. 'An Expert System for Oncology Protocol Management'. In Bruce G. Buchanan and Edward H. Shortliffe. *Rule Based Expert Systems*. Reading, MA: Addison-Wesley, 1984, 653-668.

Simmons, Henry and Paul Stolley. 'This is Medical Progress? Trends and Consequences of Antibiotic Use in the United States'. *Journal of the American Medical Association* 227 (1974): 1023-1028.

Spellberg, Brad. 'Antibiotic Judo: Working Gently with Prescriber Psychology to Overcome Inappropriate Use'. *JAMA Internal Medicine* 174, no. 3 (2014): 432-433.

Srole, Leo et al. *Mental Health in the Metropolis: The Midtown Manhattan Study*. New York: McGraw-Hill, 1962.

Tanimoto, Taffee. 'IBM Type 704 Medical Diagnosis Program'. *IRE Transactions on Medical Eletronics* ME-7 (1960): 280-283.

Taulbee, E.C., and B.D. Sisson. 'Configurational Analysis of MMPI Profiles of Psychiatric Groups'. *Journal of Consulting Psychology* 21 (1957): 413-417.

Tonkin-Crine, Sarah, Ann Sarah Wlaker, Chris C. Butler. ‘Contribution of Behavioural Science to Antibiotic Stewardship’. *British Medical Journal* 350 (2015): h3413.

Topol, Eric. *The Creative Destruction of Medicine: How the Digital Revolution Will Create Better Health Care.* New York: Basic Books, 2012.

------ *Deep Medicine: How Artificial Intelligence Can Make Health Care Human Again.* New York: Basic Books, 2019.

------ *The Patient Will See You Now: The Future of Medicine Is in Your Hands.* New York: Basic Books, 2015.

Toronto, Alan F., L. George Veasy, and Homer R. Warner. ‘Evaluation of a Computer Program for Diagnosis of Congential Heart Disease’. *Progess in Cardioveascular Disease* 5, no. 4 (1963): 362-377.

Tversky, Amos and Daniel Kahneman. ‘Availability: A Heuristic for Judging Frequency and Probability’. *Cognitive Psychology* 5, no. 2 (1973): 207-232.

------ ‘Judgement under Uncertainty: Heuristics and Biases’. *Science* 185, no. 4157 (1974): 1124-1131.

Vailhinger, Hans. *The Philosophy of ‘As if’: A System of the Theoretical, Practical and Religious Fictions of Mankind*, trans. C. K. Ogden. New York: Harcourt Brace, 1925.

van Melle, William. ‘The Strcuture of the MYCIN System’. In Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley Publishing Company, 1984, 67-77.

van Woerkom, Adrianus. ‘Program for a Diagnostic Model’. *IRE Transactions on Medical Electronics* ME-7 (1960): 220.

van Woerkom, Adrianus and Keeve Brodman. ‘Statistics for a Diagnostic Model’. *Biometrics* 17 (1961): 299-318.

von Neumann, John and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton: Princeton Univrsity Press, 1944.

Verghese, Abraham. ‘Culture Shock—Patient as Icon, Icon as Patient’. *New England Journal of Medicine* 359, no. 26 (2008): 2748-2751

Verghese, Abraham, Nigam H. Shah, and Robert A. Harrington. ‘What This Computer Needs Is a Physician: Humanism and Artificial Intelligence’. *JAMA* 319, no. 1 (2018): 19-20.

Vranas, K.C., J.K. Jopling, T.E. Sweeny, M.C. Ramsey, A.S. Milstein, C.G. Slatore, G.J. Escobar, V.X. Liu. ‘Identifying Distinct Subgroups of ICU Patients: A Machine Learning Approach’. *Critical Care Medicine* 45, no. 10 (2017): 1607-1615.

Warner, Homer R., Alan F. Toronto, and L. George Veasey. ‘Experience with Bayes’ Theorem for Computer Diagnosis of Congenital Heart Disease’. *Annals of the New York Academy of Science* 115 (1964): 558-567.

Warner, Homer R., Alan F. Toronto, L. George Veasey, and Robert Stephenson. ‘A Mathematical Approach to Medical Diagnosis: Application to Congential Heart Disease’. *Journal fo the American Medical Association* 177, no. 3 (1961): 177-183.

Weider, Arthur, Keeve Brodman, Bela Mittelmann, David Wechsler, and Harold Wolff. ‘The Cornell Index: A Method for Quickly Assaying Personality and Psychosomatic Disturbances, to Be Used as an Adjunct to Interview’. *Psychosomatic Medicine* 8 (1946): 411–413.

------ ‘Cornell Service Index: A Method for Quickly Assaying Personality and Psychosomatic Disturbances in Men in the Armed Forces’. *War Medicine* 7 (1945): 209–213.

Weider, Arthur, Bela Mittelmann, David Wechsler, Harold Wolff, and Margaret Meixner. ‘The Cornell Selectee Index: A Method for Quick Testing of Selectees for the Armed Forces’. *Journal of the American Medical Association* 124 (1944): 224–228.

Weiss, Sholom, Casimir Kulikowski, Saul Amarel, and Aran Safir. ‘A Model-Based Method for Computer-Aided Medical Decision-Making’. *Artificial Intelligence* 11 (1978): 145-172.

Wennberg, John and Alan Gittelsohn. ‘Small Area Variations in Health Care Delivery’. *Science* 182, no. 4117 (1973): 1102-1108.

Wiener, Norbert. *Cybernetics: Or, Control and Communication in the Animal and the Machine*, New York: J. Wiley, 1948.

Williams, Brian. ‘Personal View’. *British Medical Journal* 2, no. 23 (June 1973): 712.

Williams, John D. *The Compleat Strategyst: Being a Primer on the Theory of Games of Strategy*. New York: McGraw-Hill, 1966.

Woodbury, Max A. ‘Inapplicabilities of Bayes’ Theorem to Diagnosis’. In C. C. Thomas. *Proceedings of the Fifth International Conference on Medical Electronics*. Springfield, IL, 1963, 860-868.

Yarnall, Stephen R. and Richard A. Kronmal. ‘Computers Aids to Medical Diagnosis—Problems and Progress’. In *Proc. Conf. Assoc. Comp. Med.* New York: ACM, 1966, 269-274.

Yu, Victor L., Bruce G. Buchanan, Edward H. Shortliffe, Sharon M. Wraith, Randall David, A. Carlisle Scott, and Stanley N. Cohen. ‘Evaluating the Performance of a Computer-Based Consultant’. *Computer Programs in Biomedicine* 9 (1979): 95-102.

Yu, Victor L., Lawrence M. Fagan, Sharon M. Wraith, William J. Clancey, A. Carlisle Scott. John Hannigan, Robert L. Blum, Bruce G. Buchanan, and Stanley N. Cohen. 'Antimicrobial Selection by a Computer: A Blinded Evaluation by Infectious Disease Experts'. *Journal of the American Medical Association* 242, no. 12 (1979): 1279-1282.

Zworykin, Vladimir. 'A Mechanized Matching Procedure for Computer Aided Differential Diagnosis'. *Medical Electronics and Biological Engineering* 1, no. 1 (1963): 85-89.

IV. Printed Secondary Sources

Abbate, Janet. *Recoding Gender: Women's Changing Participating in Computing*. Cambridge, MA: MIT Press, 2017.

Abraham, Tara H. *Rebel Genius: Warren S. McCulloch's Transdisciplinary Life in Science*. Cambridge, MA: MIT Press, 2016.

Abramson, Albert. *The History of Television, 1880-1941*. Jefferson, NC: McFarland, 1987.

Agar, Jon. What Difference Did Computers Make? *Social Studies of Science* 36, no. 6 (2006): 869-907.

Akera, Atsushi. *Calculating a Natural World: Scientists, Engineers, and Computers During the Rise of U.S. Cold War Research*. Cambridge, MA: MIT Press, 2008.

Akrich, Madeline. 'The De-Scription of Technical Objects'. In Wiebe E. Bijker and John Law (eds.). *Shaping Technology / Building Society: Studies in Sociotechnical Change*. Cambridge, MA: MIT Press 1992), 205-224.

Anderson, Margo. *The American Census: A Social History*. New Haven: Yale University Press, 1988.

Anderson, Warwick. 'The Reasoning of the Strongest: The Polemics of Skill and Science in Medical Diagnosis'. *Social Studies of Science* 22, no. 4 (1992): 653-684.

Aronova, Elena, Christine von Oertzen, and David Sepkoski. 'Introduction: Historicizing Big Data'. *Osiris* 32 (2017): 1-17.

Aronowitz, Robert. 'The Converged Experience of Risk and Disease'. *The Milbank Quarterly* 87, no. 2 (2009): 417-442.

------ *Making Sense of Illness: Science, Society, and Disease*. Cambridge and New York: Cambridge University Press, 1998.

------ *Risky Medicine: Our Quest to Cure Fear and Uncertainty*. Chicago and London: University of Chicago Press, 2015.

Baron, R. J. ‘Quality Improvement with an Electronic Health Record: Achievable, but Not Automatic’. *Annals of Internal Medicine* 147 (2007): 548–552.

Berg, Marc. *Rationalizing Medical Work: Decision Support Techniques and Medical Practices*. Cambridge, MA: MIT Press, 1997.

------ ‘Turning a Practice into a Science: Reconceptualizing Postwar Medical Practice’. *Social Studies of Science* 25, no. 3 (1995): 437-476.

Berg, Marc and Geoffrey Bowker. ‘The Multiple Bodies of the Medical Record: Toward a Sociology of an Artifact’. *The Sociological Quarterly* 38, no. 3 (1997): 513-537.

Berg, Marc, Klasien Horstman, Saskia Plass, and Michelle van Heusden. ‘Guidelines, Professionals and the Production of Objectivity: Standardization and the Professionalism of Insurance Medicine’. *Sociology of Health and Illness* 22, no. 6 (2000): 765-791.

Berg, Marc and Annemarie Mol (eds.). *Differences in Medicine: Unraveling Practices, Techniques, and Bodies*. Durham and London: Duke University Press, 1998.

Blair, Anne. *Too Much to Know: Managing Scholarly Information before the Modern Age*. New Haven: Yale University Press, 2010.

Bleich, Howard. ‘Keeve Brodman and the Cornell Medical Index’. *MD Computing* 13, no. 2 (1996): 119-120, 122, 124.

Blum, B. and K. Duncan (eds.). *A History of Medical Informatics: Proceedings of the Association for Computing Machinery Conference on History of Medical Informatics*. New York: ACM Press, 1990.

Bothwell, Laura E., Jeremy A. Greene, Scott H. Podolsky, and David S. Jones. ‘Assessing the Gold Standard—Lessons from the History of RCTs’. *New England Journal of Medicine* 374 (2016): 2175-2181.

Bouk, Dan. *How Our Days Became Numbered: Risk and the Rise of the Statistical Individual*. Chicago: University of Chicago Press, 2015.

Brandt, Allan. *No Magic Bullet: A Social History of Venereal Disease in the United States Since 1880*. Oxford: Oxford University Press, 1987.

Brown, Nik and Andrew Webster. *New Medical Technologies and Society: Reordering Life*. Cambridge and Maldan: Polity, 2004.

Buchanan, Roderick D. ‘The Development of the Minnesota Multiphastic Personality Inventory’. *Journal of the History of the Behavioral Science* 30 (1994): 148-161.

------ ‘Ink Blots or Profile Plots: The Rorschach versus the MMPI as the Right Tool for a Science-Based Profession’. *Science, Technology, & Human Values* 22, no. 2 (1997): 168-206.

Buolamwini, Joy, and Timnit Gebru. ‘Gender Shades: Intersectioanl Accuracy Disparities in Commercial Gender Classification’. *Proceedings of Machine Learning Research* 81 (2018): 1-15.

Campbell-Kelly, Martin and William Aspray. *Computer: A History of the Information* Machine. New York: Basic Books, 1996.

Carpenter, Daniel P. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton: Princeton University Press, 2010.

Carroll, Lewis. *Through the Looking-Glass*. Mineola, NY: Drover Publications Inc., 1990.

Cartwright, Lisa. *Screening the Body: Tracing Medicine’s Visual Culture*. Minneapolis: University of Minnesota Press, 1995.

Ceruzzi, Paul. ‘When Computers Were Human’. *IEEE Annals of the History of Computing* 13, no. 3 (1991): 237-244.

Choi, Katherine, Yevgeniy Gitelman, and David A. Asch. ‘Subscribing to Your Patients—Reimagining the Future of Electronic Health Records’. *New England Journal of Medicine* 378 (2018): 1960–1962.

Cohen-Cole, Jamie. *The Open Mind: Cold War Politics and the Sciences of Human Nature*. Chicago: University of Chicago, 2014.

Collen, Morris (ed.). *A History of Medical Informatics in the United States, 1950 to 1990*. Indianapolis, IN: American Medical Informatics Association, 1995.

Collen, Morris and Marion Ball (eds.). *The History of Medical Informatics in the United States*, 2nd edition. London: Springer, 2015.

Collins, Harry. *Artificial Experts: Social Knowledge and Intelligence Machines*. Cambridge, MA: MIT Press, 1992.

Collins, Harry and Trevor Pinch. *Dr. Golem: How to Think about Medicine*. Chicago: University of Chicago Press, 2005.

Coon, Deborah J. ‘Standardizing the Subject: Experimental Psychologists, Introspection, and the Quest for a Technoscientific Ideal’. *Technology and Culture* 34, no. 4 (1993): 757-783.

Copeland, B. Jack. ‘Colossus: Its Origins and Originators’. *IEEE Annals of the History of Computing* 26, no. 4 (2004): 38-45.

Crenner, Christopher. ‘Diagnosis and Authority in the Early Twentieth-Century Medical Practice of Richard C. Cabot’. *Bulletin of the History of Medicine* 76, no. 1 (2002): 30-55.

------ *Private Practice: In the Early Twentieth-Century Medical Office of Dr. Richard Cabot*. Baltimore: Johns Hopkins University Press, 2005.

Crowther-Heyck, Hunter. *Age of System: Understanding the Development of Modern Social Science*. Baltimore: Johns Hopkins University Press, 2015.

------ *Herbert A. Simon: The Bounds of Reason in Modern America*. Baltimore: Johns Hopkins University Press, 2005.

Dallek, Robert. *Flawed Giant: Lyndon Johnson and His Times, 1961-1973*. Oxford: Oxford University Press, 1998.

Daston, Lorraine. *Classical Probability in the Enlightenment*. Princeton: Princeton University Press, 1988.

------ *Science in the Archive: Pasts, Presents, Futures*. Chicago and London: University of Chicago Press, 2017.

Davis, Audrey. 'Life Insurance and the Physical Examination: A Chapter in the Rise of American Medical Technology'. *Bulletin of the History of Medicine* 55 (1981): 392-406.

Deaton, Angus and Nancy Cartwright. 'Understanding and Misunderstanding Randomized Controlled Trials'. *Social Science & Medicine* 210 (2018): 2-21.

Dick, Stephanie. 'AfterMath: The Work of Proof in the Age of Human-Machine Collaboration'. *Isis* 102, no. 3 (2011): 494-505.

------ 'Artificial Intelligence'. *Harvard Data Science Review* 1 (2019).

------ 'Of Models and Machines: Implementing Bounded Rationality'. *Isis* 106, no. 3 (2015): 623-634.

Dolan, Brian and Allison Tillack. 'Pixels, Patterns and Problems of Vision: The Adaptation of Computer-Aided Diagnosis for Mammography in Radiological Practice in the U.S.'. *History of Science* 48, no. 2 (2010): 227-249.

Dressel, Julia, and Hany Farid. 'The accuracy, fairness, and limits of predicting recidivism'. *Science Advances* 4, no. 1 (2018): eaao5580.

Edgerton, David. *The Shock of the Old: Technology and Global History Since 1900*. Oxford and New York: Oxford University Press, 2007.

Edwards, Paul. *The Closed World: Computers and the Politics of Discourse in Cold War America.* Cambridge, MA: MIT Press, 1996.

------ *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming.* Cambridge, MA: MIT Press, 2010.

Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: Macmillan, 2017.

Faulkner, Alex. *Medical Technology into Healthcare and Society: A Sociology of Devices, Innovation and Governance*. New York and London: Palgrave Macmillan, 2009.

Feigenbaum, Edward and Bruce Buchanan. 'DENDRAL and Meta-DENDRAL: Roots of Knowledge Systems and Expert System Applications'. *Artificial Intelligence* 59 (1993): 233-240.

Fienberg, Stephen E. 'When Did Bayesian Inference Become "Bayesian"?'. *Bayesian Analysis* 1, no. 1 (2006): 1-40.

Forsythe, Diana. 'Engineering Knowledge: The Construction of Knowledge in Artificial Intelligence'. *Social Studies of Science* 23, no. 3 (1993): 445-477.

------ *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford: Stanford University Press, 2002.

Fox, Daniel M. 'From Piety to Platitudes to Pork: The Changing Politics of Health Workforce Policy'. *Journal of Health Politics, Policy and Law* 21 (1996): 825-844.

------ *Health Policies, Health Politics: The British and American Experience, 1911-1965*. Princeton: Princeton University Press, 1986.

Fox, Rene. 'The Evolution of Medical Uncertainty'. *The Milbank Quarterly* 58, no. 1 (1980): 1-49.

------ 'Training for Uncertainty'. In *The Student Physician*. Cambridge, MA: Harvard University Press, 1957.

Frängsmyr, Tore, J. L. Heilbron, and Robin Rider, eds. *The Quantifying Spirit in the 18th Century*. Berkeley: University California Press, 1990.

Fuller, Matthew (ed.). *Software Studies: A Lexicon.* Cambridge, MA: MIT Press, 2008.

Galison, Peter. 'The Ontology of the Enemy: Norbert Wiener and the Cybernetic Vision'. *Critical Inquiry* 21 (1994): 228-266.

Gillmor, C. Stewart. *Fred Terman at Stanford: Building a Discipline, a University, and Silicon Valley*. Stanford: Stanford University Press, 2004.

Gitelman, Lisa. *Always Already New: Media, History, and the Data of Culture*. Cambridge, MA: MIT Press, 2006.

------ *Paper Knowledge: Toward. Media History of Documents*. Durham: Duke University Press, 2014.

------ *"Raw Data" Is an Oxymoron*. Cambridge, MA: MIT Press, 2013.

Gong, Bo, James P. Nugent, William Guest, William Parker, Paul J. Chang, Faisal Khosa, Savvas Nicolaou. 'Influence of Artificial Intelligence on Canadian Medical Students' Preference for Radiology Specialty: A National Survey Study'. *Academic Radiology* 26, no. 4 (2019): 566-577.

Gordin, Michael. *Scientific Babel: How Science Was Done Before and After Global English*. London and Chicago: University of Chicago Press, 2015).

Grames, Jr., Martin. *Science Fiction Theatre: A History of the Television Program, 1955-1957*. Albany, GA: BearManor Media, 2011.

Greene, Jeremy A. *Generic: The Unbranding of Modern Medicine*. Baltimore: Johns Hopkins University Press, 2014.

------ *Prescribing by Numbers: Drugs and the Definition of Disease*. Baltimore: Johns Hopkins University Press, 2007.

------ 'When Television Was a Medical Device'. *Humanities: The Magazine for the National Endowment for the Humanities* 38, no. 2 (2017).

Greene, Jeremy and Andrew Lea. 'Digital Futures Past—The Long Arc of Big Data in Medicine'. *New England Journal of Medicine* 381 (2019): 480-485.

Greene, Jeremy and Scott Podolsky. 'Reform, Regulation, and Pharmaceuticals—The Kefauver-Harris Amendments at 50'. *New England Journal of Medicine* 367, no. 16 (2012): 1481-1483.

Grove, William. 'Clinical versus Statistical Prediction: The Contribution of Paul E. Meehl'. *Journal of Clinical Psychology* 61, no. 10 (2005): 1233-1243.

Guenther, Katja and Volker Hess. 'Soul Catchers: The Material Culture of the Mind Sciences'. *Medical History* 60, no. 3 (2016): 301-307.

Harris, Nancy Lee and Robert Sculley. 'The, Clinicopathological Conferences (CPCs)'. In *Keen Minds to Explore the Dark Continents of Disease*. Boston: Massachusetts General Hospital and Harvard Medical School, 2011.

Heath, Christian, Paul Luff, and Marcus Sanchez Svensson. 'Technological and Medical Practice'. *Sociology of Health and Illness* 25, no. 3 (2003): 75-96.

Heims, Steve J. *The Cybernetics Group*. Cambridge, MA: MIT Press, 1991.

Hepler-Smith, Evan. '"A Way of Thinking Backwards": Computing and Method in Synthetic Organic Chemistry'. *Historical studies in the Natural Sciences* 48, no. 3 (2018): 300-337.

Hess, Volker. 'A Paper Machine of Clinical Research in the Early Twentieth Century'. *Isis* 109, no. 3 (2018); 473-493.

Hess, Volker and J. Andrew Mendelsohn. 'Case and Series: Medical Knowledge and Paper Technology, 1600-1900'. *History of Science* 48, no. 3-4 (2010): 287-314.

Hofmann, Bjørn. 'The Technological Invention of Disease'. *Medical Humanities* 27, no. 1 (2001): 10-19.

Hughes, Agatha and Thomas Parke Hughes. *Systems, Experts, and Computers: The Systems Approach in Management and Engineering, World War II and After.* Cambridge, MA: MIT Press, 2000.

Igo, Sarah. *The Averaged American: Surveys, Citizens, and the Making of a Mass Public*. Cambridge: Harvard University Press, 2007.

------ *The Known Citizen: A History of Privacy in Modern America*. Cambridge, MA: Harvard University Press, 2018.

Johnson, Ericka and Boel Berner. *Technology and Medical Practice: Blood, Guts, and Machines*. Farnham and Burlington: Ashgate Publishing Company, 2010.

Jones, David S. *Broken Hearts: The Tangled History of Cardiac Care*. Baltimore: Johns Hopkins University Press, 2013.

------ 'Visions of a Cure: Visualization, Clinical Trials, and Controversies in Cardiac Therapeutics, 1968-1998'. *Isis* 91, no. 3 (2000): 504-541.

Jones, Matthew L. 'How We Became Instrumentalists (Again): Data Positivism since World War II'. *Historical Studies in the Natural Sciences* 48, no. 5 (2018): 673-684.

Kay, Lily. *Who Wrote the Books of Life? A History of the Genetic Code*. Stanford: Stanford University Press, 2000.

Kaplan, Bonnie. 'The Computer Prescription: Medical Computing, Public Policy, and Views of History'. *Science, Technology, & Human Values* 20, no. 1 (1995): 5-38.

Keating, Peter and Alberto Cambrosio. *Biomedical Platforms: Realigning the Normal and the Pathological in Late-Twentieth-Century Medicine*. Cambridge, MA: MIT Press, 2003.

------ 'Cancer Clinical Trials: The Emergence and Development of a New Style of Practice'. *Bulletin of the History of Medicine* 81 (2007): 197-223.

Klein, Judy L., Lorraine Daston, Rebecca Lemov, Thomas Sturm, and Michael D. Gordin. *How Reason Almost Lost Its Mind: The Strange Career of Cold War Rationality*. Chicago and London: University of Chicago Press, 2013.

Kline, Ronald. *The Cybernetics Moment: Or, Why We Call Our Age the Information Age*. Baltimore: Johns Hopkins University Press, 2015.

Kuperman, Gilad J., Reed M. Gardner, and T. Allan Pryor. *HELP: A Dynamic Hospital Information System*. New York: Springer-Verlag, 1991.

Kurzweil, Raymond. *The Age of Intelligent Machines.* Cambridge, MA: MIT Press, 1996.

Latour, Bruno. *Reassembling the Social: An Introduction to Actor-Network Theory.* Oxford: Oxford University Press, 2005.

Latour, Bruno and Stephen Woolgar. *Laboratory Life: The Construction of Scientific Facts.* Beverely Hills: Sage, 1979.

Lawrence, Christopher. Incommunicable Knowledge: Science, Technology and the Clinical Art in Britain, 1850-1914'. *Journal of Contemporary History* 20, no. 4 (1985): 503-520.

------ 'Still Incommunicable: Clinical Holists and Medical Knowledge in Interwar Britain'. In *Greater Than the Parts: Holism in Biomedicine, 1920-1950*. Oxford: Oxford University Press, 1998.

Lea, Andrew. 'Computerizing Diagnosis: Keeve Brodman and the Medical Data Screen'. *Isis* 110, no. 2 (2019): 228-249.

Ledley, Robert S. 'Medical Informatics: A Personal View of Sowing the Seeds'. In *A History of Medical Informatics*. New York: ACM Press, 1990.

Lehoux, Pascale. *The Problem of Health Technology: Policy Implications for Modern Health Care Systems*. New York and London: Routledge, 2006.

Lemov, Rebecca. *Database of Dreams: The Lost Quest to Catalog Humanity*. Hew Haven: Yale University Press, 2015.

Leslie, Stuart W. 'Playing the Education Game to Win: The Military and Interdisciplinary Research at Stanford'. *Historical Studies in the Physical and Biological Sciences* 18, no. 1 (1987): 55-88.

Lowen, Rebecca S. *Creating the Cold War University: The Transformation of Stanford*. Berkeley: University of California Press, 1997.

------ '"Exploiting a Wonderful Opportunity": The Patronage of Scientific Research at Stanford University, 1937-1965'. *Minerva* 30, no. 3 (1992): 391-421.

Löwy, Ilana. 'La standardization de l'inconnu: Les protocols thérapeutiques en cancérologie'. *Techniques and Culture* 25/26 (1995): 73–108.

Ludmerer, Kenneth. *Learning to Heal: The Development of American Medical Education*. New York: Basic Books, 1985.

------ *Time to Heal: American Medical Education from the Turn of the Century to the Era of Managed Care*. Oxford: Oxford University Press, 1999.

MacKenzie, Donald. *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press, 2006.

Mahoney, Michael. 'The Histories of Computing(s)'. *Interdisciplinary Science Reviews* 30, no. 2 (2005): 119-135.

Magoun, Alexander. *Television: The Life Story of a Technology*. Westport, CT: Greenwood Press, 2007.

Marks, Harry M. 'Medical Technologies: Social Contexts and Consequences'. In W. F. Bynum and Roy Porter. *Companion Encyclopedia of the History of Medicine*, Vol. 2. New York: Routledge, 1993, 1592-1618.

------ *The Progress of Experiment: Science and Therapeutic Reform in the United States, 1900-1990*. Cambridge: Cambridge University Press, 1997.

McCarthy, John. 'History of Lisp'. In *History of Programming Languages I.* New York, 1978.

McCray, W. Patrick. 'How Astronomers Digitized the Sky'. *Technology and Culture* 55 (2014): 908-944.

------ *The Visioneers: How a Group of Elite Scientists Pursued Space Colonies, Nanotechnologies, and a Limitless Future*. Princeton: Princeton University Press, 2013.

McGrayne, Sharon. *The Theory that Would Not Die*. New Haven: Yale University Press, 2011.

Meehl, Paul E. 'Paul E. Meehl'. In G. Lindzey. *A History of Psychology in Autobiography*. Stanford: Stanford University Press, 1989, 337-389.

Mendelsohn, J. Andrew. 'Empiricism in the Library: Medicine's Case Histories'. In *Science in the Archives: Pasts, Presents, Futures*. Chicago: University of Chicago Press, 2017.

------ 'The World on a Page: Making a General Observation in the Eighteenth Century'. In *Histories of Scientific Observation*. Chicago: University of Chicago Press, 2011.

Mody, Cyrus C.M. *The Long Arm of Moore's Law: Microelectronics and American Science*. Cambridge, MA: MIT Press, 2017.

Mol, Annemarie. *The Body Multiple: Ontology in Medical Practice*. Durham: Duke University Press, 2002.

------ *The Logic of Care: Health and the Problem of Patient Choice*. London and New York: Routledge, 2008.

Mulkay, Michael, Trevor Pinch, and Malcolm Ashmore. 'Colonizing the Mind: Dilemmas in the Application of Social Science'. *Social Studies of Science* 17, no. 2 (1987): 231-256.

Mullainathan, Sendhil, and Ziad Obermeyer. ‘Does Machine Learning Automate Moral Hazard and Error?’. *American Economic Review* 107, no. 5 (2017): 476-480.

Müller, Albert and Karl Müller. *An Unfinished Revolution? Heinz von Foerster and the Biology Computer Laboratory (BCL) 1958-1976*. Vienna: Edition Echoraum, 2007.

Nettleton, Sarah. ‘The Emergence of E-Scaped Medicine?’. *Sociology* 38, no. 4 (2004): 662-679.

Nilsson, Nils. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge and New York: Cambridge University Press, 2010.

Noble, Safiya. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.

November, Joseph. *Biomedical Computing: Digitizing Life in the United States.* Baltimore: Johns Hopkins University Press, 2012.

------ ‘Early Biomedical Computing and the Roots of Evidence-Based Medicine’. *IEEE Annals of the History of Computing* 33, no. 2 (2011): 9-23.

Oldfield, Benjamin J. and David S. Jones. ‘Languages of the Heart: The Biomedical and the Metaphorical in American Fiction’. *Perspectives in Biology and Medicine* 57, no. 3 (2014): 424-442.

Olin-Lauritzen, Sonja and Lars-Christer Hyden. *Medical Technologies and the Life World: The Social Construction of Normality*. London and New York: Routledge, 2007.

O’Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown, 2016.

Ornstein, Severo. *Computing in the Middle Ages: A View from the Trenches, 1955-1983*. Bloomington: 1stBooks, 2002.

Ostherr, Kirsten. *Medical Visions: Producing the Patient through Film, Television, and Imaging Technologies*. Oxford: Oxford University Press, 2013.

Oudshoorn, Nelly. *Telecare Technologies and the Transformation of Healthcare*. Basingstoke: Palgrave Macmillan, 2011.

Oudshoorn, Nelly and Trevor Pinch. *How Users Matter: The Co-Construction of Users and Technologies*. Cambridge, MA: MIT Press, 2003.

Pickering, Andrew. *The Cybernetic Brain: Sketches of Another Future*. Chicago: University of Chicago Press, 2010.

Plasek, Aaron. ‘On the Cruelty of Really Writing a History of Machine Learning’. *IEEE Annals of the History of Computing* 38, no. 4 (2016): 6-8.

Podolsky, Scott H. *The Antibiotic Era: Reform, Resistance, and the Pursuit of a Rational Therapeutics*. Baltimore: Johns Hopkins University Press, 2015.

Podolsky, Scott H., Robert Bud, Christoph Gradmann, Bård Hobaek, Claas Kirchhelle, Tore Mitvedt, María Jesús Santesmases, Ulrike Thoms, Dag Berild, and Anne Kveim Lie. 'History Teaches Us That Confronting Antibiotic Resistance Requires Stronger Global Collective Action'. *Journal of Law, Medicine and Ethics* 43 (2015): 27-32.

Polachek, Harry. 'Before the ENAIC'. *IEEE Annals of the History of Computing* 19, no. 2 (1997): 25-30.

Price, W. N. 'Regulating Black-box Medicine'. *Michigan Law Review* 116, no. 1 (2017): 421–474.

Radin, Joanna. '"Digital Natives": How Medical and Indigenous Histories Matter for Big Data'. *Osiris* 32 (2017): 43-64.

Reiser, Stanley Joel. *Medicine and the Reign of Technology*. Cambridge and New York: Cambridge University Press, 1978.

------ *Technological Medicine: The Changing World of Doctors and Patients*. Cambridge and New York: Cambridge University Press, 2009.

Rosenberg, Charles. 'Framing Disease: Illness, Society, and History'. In *Framing Disease*. New Brunswick: Rutgers University Press, 1992.

------ 'The Tyranny of Diagnosis: Specific Entities and Individual Experience'. *Milbank Quarterly* 80, no. 2 (2002): 237-260.

Russell, Andrew L. and Lee Vinsel. 'After Innovation, Turn to Maintenance'. *Technology and Culture* 59, no. 1 (2018): 1-25.

------ 'Making Maintainers: Engineering Education and the Ethics of Care'. In *Does American Need More Innovators?* Cambridge, MA: MIT Press, 2019.

Schilling, Rebecca and Stephen T. Casper. 'Of Psychometric Means: Starke R. Hathaway and the Popularization of the Minnesota Multiphasic Personality Inventory'. *Science in Context* 28, no. 1 (2015): 77-98.

Schlich, Thomas and Christopher Crenner (eds.). *Technological Change in Modern Surgery: Historical Perspectives on Innovation*. Melton: University of Rochester Press, 2017.

Seising, Rudolf. 'Robert S. Ledley, 1926-2012'. *Artificial Intelligence in Medicine* 57, no. 1 (2013): 1-7.

Sepkoski, David. 'The Database before the Computer?'. *Osiris* 32 (2017): 175-201.

Shapin, Steven. ‘Pump and Circumstance: Robert Boyle’s Literary Technology’. *Social Studies of Science* 14, no. 4 (1984): 481-520.

------ ‘A Taste of Science: Making the Subjective Objective in the California Wine World’. *Social Studies of Science* 46, no. 3 (2016): 436-460.

Shapin, Steven and Simon Schaffer. *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton: Princeton University Press, 1985.

Shustek, Len. ‘An Interview with Ed Feigenbaum’. *Communications of the ACM* 53, no. 6 (2010): 41-45.

Sittig, Dean, Joan Ash, and Robert Ledley. ‘The Story Behind the Development of the First Whole-body Computerized Tomography Scanner as Told by Robert S. Ledley’. *Journal of the American Medical Informatics Association* 13, no. 5 (2006): 465-469.

Sontag, Susan. *AIDS and Its Metaphors.* New York: Farrar, Straus and Giroux, 1989.

------ *Illness as Metaphor*. New York: Vintage Books, 1979.

Stanton, Jennifer. ‘Making Sense of Technologies in Medicine’. *Social History of Medicine* 12 (1999): 437-448.

Starr, Paul. *The Social Transformation of American Medicine: The Rise of a Sovereign Profession and the Making of a Vast Industry*. New York: Basic Books, 1982.

Steingart, Alma. ‘A Group Theory of Group Theory: Collaborative Mathematics and the “Uninvention” of a 1000-Page Proof’. *Social Studies of Science* 42, no. 2 (2012): 185-213.

Stevens, Hallam. ‘A Feeling for the Algorithm: Working Knowledge and Big Data in Biology’. *Osiris* 32 (2017): 151-174.

------ *Life out of Sequence: A Data-Driven History of Bioinformatics*. Chicago and London: University of Chicago Press, 2013.

Strasser, Bruno J. and Paul Edwards. ‘Big Data Is the Answer … But What is the Question?’. *Osiris* 32 (2017): 328-345.

Strasser, Bruno J. *Collecting Experiments: Making Big Data Biology*. Chicago: University of Chicago Press, 2019.

------ ‘The Experimenter’s Museum: GenBank, Natural History, and the Moral Economies of Biomedicine’. *Isis* 102, no. 1 (2011): 60-96.

Strasser, Susan. *Never Done: A History of American Housework*. New York: Pantheon, 1982.

Strickland, Stephen P. *The History of the Regional Medical Programs: The Life and Death of a Small Initiative of the Great Society*. Lanham, MD: University Press of America, 2000.

Suchman, Lucy. *Human-Machine Reconfigurations*. Cambridge: Cambridge University Press, 2006.

Tillack, Allison A. and Richard S. Breiman. 'Renegotiating Expertise: An Examination of PACS and the Challenges to Radiology Using a Medical Anthropologic Aproach'. *Journal of the American College of Radiology* 9, no. 1 (2012): 64-68.

Timmermann, Carsten and Julie Anderson (eds.). *Devices and Designs: Medical Technologies in Historical Perspective*. New York: Palgrave Macmillan, 2006.

Timmermans, Stefan and Marc Berg. *The Gold Standard: The Challenge of Evidence-Based Medicine and Standardization in Health Care*. Philadelphia: Temple University Press, 2003.

------ 'The Practice of Medical Technology'. *Sociology of Health and Illness* 25 (2003): 97-114.

------ 'Standardization in Action: Achieving Local Universality through Medical Protocols'. *Social Studies of Science* 27, no. 2 (1997): 273-305.

Tobbell, Dominique. *Pills, Power, and Policy: The Struggle for Drug Reform in Cold War America and Its Consequences*. Berkeley: University of California Press, 2011.

Vayena, Effy, Alessandro Blasimme, I. Glenn Cohen. 'Machine Learning in Medicine: Addressing Ethical Challenges'. *PLOS Medicine* 15, no. 11 (2018): e1002689.

von Oertzen, Christine. 'Machineries of Data Power: Manual versus Mechanical Census Compilation in Nineteenth-Century Europe'. *Osiris* 32 (2017): 129-150.

Wailoo, Keith. *Drawing Blood: Technology and Disease Identity in Twentieth-Century America*. Baltimore: Johns Hopkins University Press, 1997.

------ *Dying in the City of the Blues: Sickle Cell Anemia and the Politics of Race and Health*. Chapel Hill: University of North Carolina Press, 2001.

Warner, John Harley. 'The Aesthetic Grounding of Modern Medicine'. *Bulletin of the History of Medicine* 88, no. 1 (2014): 1-47.

------ *The Therapeutic Perspective: Medical Practice, Knowledge, and Identity*. Cambridge, MA: Harvard University Press, 1986.

Weisz, George, Alberto Cambrosio, Peter Keating, Loes Knaapen, Thomas Schlich, and Virginie J. Tournay. 'The Emergence of Clinical Practice Guidelines'. *The Milbank Quarterly*. 85, no. 4 (2007): 691-727.

Whitfield, Nicholas and Thomas Schlich. ‘Skills through History’. *Medical History* 59, no. 3 (2015): 349-360.

Wolfe, Audra J. ‘Germs in Space: Joshua Lederberg, Exobiology, and the Public Imagination, 1958-1964’. *Isis* 93, no. 2 (2002): 183-205.

Yu, Kun-Hsing, Andrew L. Beam, and Isaac S. Kohane. ‘Artificial Intelligence in Healthcare’. *Nature Biomedical Engineering* 2 (2018): 719-731.

V. Oral Histories

Barnett, Octo. ‘From “Farm Boy” to Director of the Laboratory of Computer Science: 2004 Interview of Octo Barnett’, interviewed by Joan S. Ash and Dean F. Sittig. In Rebecca M. Goodwin, Joan S. Ash, and Dean F. Sittig (eds.). *Conversations with Medical Informatics Pioneers: An Oral History Collection*. Bethesda: U.S. National Library of Medicine, 2015.

Cohen, Stanley N. ‘Science, Biotechnology, and Recombinant DNA: A Personal History’, oral history conducted by Sally Smith Hughes in 1995. Regional Oral History Office, Bancroft Library, University of California at Berkeley, 2009.

Arthur L. Norberg, ‘An Interview with Bruce G. Buchanan’. Oral History 230, 11-12 Jun. 1991. Charles Babbage Institute. University of Minnesota.

Warner, Homer R. ‘Founder of the HELP System and the Utah Medical Informatics Program’, interviewed by Dean F. Sittig. In Rebecca M. Goodwin, Joan S. Ash, and Dean F. Sittig (eds.). *Conversations with Medical Informatics Pioneers: An Oral History Collection*. Bethesda: U.S. National Library of Medicine, 2015.

VI. Unpublished Theses and Conference Papers

Dick, Stephanie A. ‘After Math: (Re)configuring Minds, Proof, and Computing in the Postwar United States’, Ph.D. Dissertation. Cambridge, MA: Harvard University, 2015.

Kaplan, Bonnie M. ‘Computers in Medicine, 1950-1980: The Relationship Between History and Policy’, Ph.D. Dissertation. Chicago, IL: University of Chicago, 1983.

November, Joseph A. ‘Digitizing Life: The Introduction of Computers to Biology and Medicine’, Ph.D. Dissertation. Princeton, NJ: Princeton University, 2006.

Greenes, Robert A. ‘A Computer-Based System for Medical Record-Keeping by Physicians’, Ph.D. Dissertation. Cambridge, MA: Harvard University, 1969.

Shortliffe, Edward H. ‘A Computer-Based Medical Record System: Approaches to the Interface Problem’, A.B. Thesis. Cambridge, MA: Harvard University, 1970.

Theodore, David. ‘Towards a New Hospital: Architecture, Medicine, and Computation, 1960-75’, Ph.D. Dissertation. Cambridge, MA: Harvard University, 2014.

Tillack, Allison A. ‘Imaging Trust: Information Technologies and the Negotiation of Radiological Expertise in the Hospital’, Ph.D. Dissertation. San Francisco, CA: University of California at San Francisco, 2012.

Warner, John Harley. ‘Vital Signs: The Transformation of the Medical Record in the Nineteenth Century’. Keynote address at ‘Media Medica: Medicine and the Challenge of New Medicine’. Conference at Center for Medical Humanities and Social Medicine. Johns Hopkins University, Baltimore, Maryland, 17 Oct. 2017.