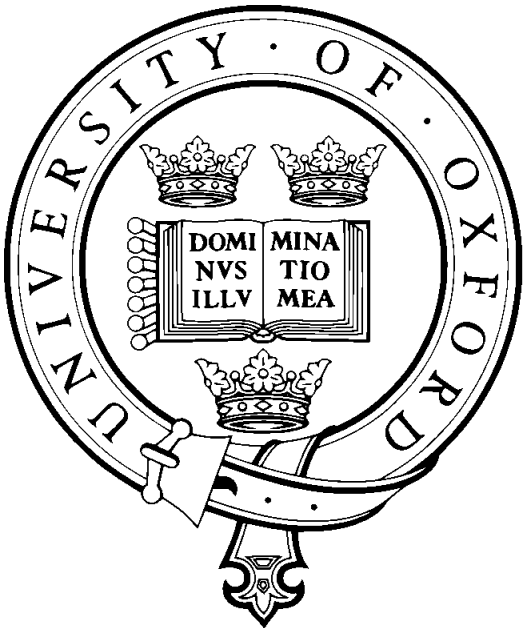
UNIVERSITY OF OXFORD

On the Role of Algorithmic Technologies in Improving Knowledge Work



Javier R. Amaya Silva
Green Templeton College
University of Oxford

Supervised by
Professor Matthias Holweg and Dr Steve New
Submitted: Hilary Term,
April 21st, 2023

This thesis is submitted to the University of Oxford, Saïd Business School, for the degree of Doctor of Philosophy

To Myriam, Ramon, and Diana

for their encouragement, care, and unconditional support along this journey

Acknowledgements

I would not have been able to conduct this research without the unconditional support of many individuals. I would like to give special thanks to my supervisors, Professor Matthias Holweg and Dr Steve New who have provided me with invaluable guidance, mentorship, and constructive feedback throughout my doctoral studies. I am also very appreciative of Dr Marc Ventresca for his all-rounded advice and Professors Thomas Lawrence and Eero Vaara's invaluable insights into academic writing as a craft. I am also grateful to Professor Mari Sako, Dr Agni Orfanoudaki, Dr Janet Smart, Dr Harvey Maylor, Dr Ken Okamura and Professor Michael Smets for their constructive comments on various iterations of this thesis. I extend my gratitude to six anonymous reviewers whose immensely helpful comments and encouragement have significantly contributed to the papers presented in this thesis. Furthermore, I thank all the members of Management Write club, as well as the participatns of the DPhil Management Presentation events at the Saïd Business School for their insightful questions and comments. I would like to acknowledge the people who have made this research possible, especially those connected to "ACME" and "Superstore", who contributed their time and allowed me to conduct the experiments and observations on their fascinating work. I am also grateful to the Saïd Foundation, and Green Templeton College for providing the resources to undertake this journey. Lastly, I would like to thank Jessica Hodgson and Anna Blas Luna for their day-to-day support throughout my doctoral journey.

Contents

Acknowledgements ..... **iii**

List of Tables ..... **viii**

List of Figures ..... **ix**

Declarations ..... **x**

Abstract ..... **xii**

Chapter 1 Introduction **1**

Chapter 2 Background and theoretical foundations **7**

2.1 Introduction 8

2.2 The concept of knowledge work 11

2.2.1 Work 12

2.2.2 Knowledge 13

2.2.3 Knowledge work 15

2.3 The concept of algorithmic technologies 17

2.3.1 A brief history of algorithmic technologies 17

2.3.2 Defining algorithmic technologies 21

2.3.3 Summary 25

2.4 Theoretical foundations 26

2.4.1 Technology and occupations 26

2.4.2 Technology, knowledge and organisational capabilities 30

2.4.3 Process Improvement 32

2.5 The interplay between algorithmic technologies and knowledge work 41

2.5.1 Research theme 1: Can algorithmic technologies improve knowledge work? 42

2.5.2 Research theme 2: How do organisations use algorithmic technologies to improve knowledge work? 44

2.5.3 Research theme 3: How do development and operations practices affect algorithmic technologies? 46

2.6 Conclusion 48

Chapter 3 Learning to (be told to) count: Implementing algorithmic recommendations in retail operations 63

3.1 Introduction 65

3.2 Theoretical foundations 68

3.2.1 Using algorithms to augment operational performance 68

3.2.2 The interplay between humans and algorithms 70

3.3 Focal problem 72

3.3.1 The study of inventory record inaccuracy 72

3.3.2 Research context 74

3.3.3 Experimental setup 77

3.4 Deep learning-based decision model 78

3.4.1 Model logic 78

3.4.2 Model's data and validation 82

3.4.3 Key variables 87

3.5 Findings 92

3.5.1 Overall performance effect 92

3.5.2 Enactment of algorithmic recommendations 95

3.5.3 Robustness of the effects 96

3.6 Discussion 98

3.6.1 Representation 99

3.6.2 Aversion 101

3.6.3 Degradation 105

3.6.4 Managerial implications 106

3.7 Conclusion 108

Chapter 4 Using algorithmic technologies to improve knowledge work ..116

4.1 Introduction 118

4.2 Theoretical background 121

4.2.1 Algorithmic technologies in improving knowledge work 121

4.2.2 A knowing-in-practice view of algorithmic technologies 124

4.2.3 Synthesis 126

4.3 Method 127

4.3.1 Research context and design 127

4.3.2 Data collection 132

4.3.3 Data analysis 137

4.4 Findings 140
4.4.1 Overcoming inertia 141
4.4.2 Transforming knowledge work 142
4.4.3 Adjusting the knowledge regime 156
4.5 Discussion 161
4.5.1 An integrated process model of adoption of algorithmic technologies in knowledge work 162
4.5.2 Technological change in knowledge work 165
4.5.3 Technology use for improving knowledge work 168
4.5.4 Managerial implications 171
4.5.5 Limitations and further research 173
4.5.6 Outlook 175
Appendix A. Interview protocol phase I 183
Appendix B. Semi-structured interview protocol phase II 184
Appendix C. Data structure 186
Chapter 5 Enacting artificial intelligence: Connecting the development of machine learning systems with their operational effectiveness 188
5.1 Introduction 190
5.2 The emergence of the MLOps development practice 193
5.2.1 Historical foundations of MLOps 193
5.2.2 Conceptual foundations of MLOps 198
5.3 Situating MLOps in organisations: A sociomaterial ontology 207
5.3.1 Sociomateriality 207
5.3.2 The sociomateriality of MLOps 212
5.4 Operational effectiveness implications of machine learning in knowledge work 217
5.4.1 How the culture of MLOps shapes knowledge production 219
5.4.2 How practices in MLOps shape knowledge outputs 224
5.4.3 How the collection of tools and artefacts of MLOps shapes knowledge outputs 227
5.5 Conclusion 230
5.5.1 Theoretical contributions: From neutrality to complexity 232
5.5.2 Managerial implications 234
Chapter 6 Conclusion 246
6.1 Theoretical contributions 247

6.2 Managerial implications 250
6.3 Limitations and further research 252

List of Tables

Table 2-1 A typology of algorithmic technologies with their defining features.....23
Table 3-1 Characteristics of the randomized store allocation for our experiment 78
Table 3-2 Drivers and variables that may affect inventory record inaccuracies. 80
Table 3-3 SKU composition of the sample data, categorized by store format 84
Table 3-4 Performance comparison for deep learning meta-models 85
Table 3-5 Examples of SKUs predicted to be in error 85
Table 3-6 Performance of the meta-model on full counting audits 86
Table 3-7 Variable definition for the impactful count metric 88
Table 3-8 Variable definitions for the sales metric 89
Table 3-9 Variable definitions for the waste metric.....90
Table 3-10 Variable definitions for sales-based availability metrics 91
Table 3-11 Overall effect of the deep learning-based model on performance 94
Table 4-1 Description of cases and data collection 136
Table 4-2 Descriptions of projects involving algorithmic technologies across cases 139
Table 4-3 Transforming work by altering specific tasks 145
Table 4-4 Transforming knowledge work by re-engineering the entire processes.....150
Table 5-1 Selected definitions of MLOps 199
Table 5-2 Examples of specific MLOps-related principles 203
Table 5-3 Examples of specific MLOps-related activities 204
Table 5-4 Examples of specific MLOps-related tools and artefacts 206

List of Figures

Figure 2-1 Conceptual model of algorithmic technologies on performance outcomes. ...10
Figure 2-2 Process improvement approaches. The main domains of application in the Standardisation and Repetition Matrix. Source: Based on Six Sigma lecture slides (Holweg) 34
Figure 2-3 Conceptual model for studying research theme 1 44
Figure 2-4 Conceptual model for understanding the transformation of knowledge work. 46
Figure 3-1 Operations analysts' decision model for correcting IRI based on periodic cycle counting routines. 75
Figure 3-2 A deep learning-based decision model for the cycle counting routine in the context of grocery and merchandise retail 79
Figure 3-3 Histogram of the absolute difference between recorded and in-hand inventory 84
Figure 3-4 Number of impactful errors detected before and after the model was deployed for traditional and algorithm-recommended counts. The shaded area represents standard errors. 93
Figure 3-5 Net effect of the deep learning-based decision model on the customer service level. 95
Figure 3-6 Comparison of counting recommendation adherence between control and treatment groups 95
Figure 3-7 Receiver operating characteristic curve for various machine learning models considered in the study. 97
Figure 3-8 The percentage of impactful inventory errors detected across all treatment stores during the trial period. 98
Figure 3-9 Store's non-adherence to counting recommendations. 103
Figure 4-1 Perceived ability to automate a task as a function of task characteristics. 143
Figure 4-2 A process model of the adoption of algorithms in knowledge work 164
Figure 5-1 Theoretical apparatus for connecting common themes of MLOps with the effects of machine learning algorithms in knowledge work through a sociomaterial lens. 212
Figure 5-2 Conceptual model connecting a machine learning design and operations practice with its effectiveness implications. 217

Declarations

The work contained within this document is original work conducted by myself with support from my supervisors Professor Matthias Holweg and Dr Steve New unless specified otherwise in the text or the following list of publications arising from this document.

Journal submissions
Using algorithmic technologies to improve knowledge work (under review – third round). *Journal of Operations Management* (Submitted September 2021)

Learning to (Be Told to) Count: Implementing Algorithmic Recommendations in Retail Operations (revising the manuscript to be resubmitted). *Journal of Operations Management*

Conference papers
How do Organisations Experience and Use Algorithms to Transform Knowledge Work? *Academy of Management Proceedings 2023.*

Using Deep Learning to Improve Inventory Record Accuracy. *POMS*, 2023.

Using Deep Learning to Improve Inventory Record Accuracy: Concept and Application. *Academy of Management Proceedings 2022.*

To Automate or not to Automate? The Impact of Algorithmic Technologies on Improving Knowledge Work. *Academy of Management Proceedings 2022.*

To automate or not to automate? The impact of algorithms on improving white-collar work. *EurOMA conference 2021*

Conference presentation
Using algorithmic technologies to improve knowledge work. *EurOMA conference 2023*

Using deep learning to improve inventory accuracy. *EurOMA conference 2022*

Do algorithms affect retail management? Evidence from a large grocery retailer *Conference on Artificial Intelligence, Machine Learning and Business Intelligence*, December 2021.

On the role of digital automation technologies on process improvement: effects on professional auditing. *Professional Service Firms Annual Conference*, July 2020

Doctoral seminars

On the role of algorithmic technologies in improving knowledge work. *EGOS Doctoral Symposium*, 2023.

On the role of digital automation technologies in process improvement. *EurOMA publishing workshop*, 2020.

Email traffic in "organised anarchies": Understanding the link between digital communication and office productivity. *EurOMA doctoral seminar*, 2019.

Abstract

This thesis explores the opportunities that novel algorithmic technologies present for improving knowledge work contexts. Operations management research, so far, has mostly discussed the effects of automation technologies in manufacturing processes, while our understanding of algorithmic technologies in knowledge work remains nascent. Addressing this gap, this thesis features three essays concerning algorithmic technologies, knowledge work and operations management. To set the stage for these essays, we begin by reviewing a set of management theories (e.g., technology and occupations, technology, knowledge and capabilities, and process improvement). Using these theories we propose an overarching framework, highlighting mechanisms and contingencies, that depicts the potential of algorithmic technologies in knowledge work. We then use this framework to motivate the three studies presented below.

The first essay, "*Learning to (Be Told to) Count: Implementing Algorithmic Recommendations in Retail Operations*", begin by exploring whether algorithmic technologies can indeed improve knowledge work tasks. To that end, we develop and implement a deep neural network decision model to address the problem of inventory record inaccuracy in retail management. Running an experiment involving 450 of the supermarkets and hypermarkets of a large grocery retailer, we estimate that our model outperforms operations analysts by 26.07% in detecting inventory record errors. The paper also shows that this improvement further boosts the retailer's performance in terms of availability, waste, and sale KPIs. We conclude by discussing the limitations of the

deep neural network decision model, implications for process improvement, and the future of retail management.

In the second essay, "*Using algorithmic technologies to improve knowledge work*" (status: third round of review in the Journal of Operations Management), we investigate the adoption of machine learning and other algorithms across four areas of a multinational energy firm. Through a multiple-case study approach, we delineate the successes and setbacks of knowledge workers as they integrate algorithms into their work, taking into account occupational and work differences. In this essay, we introduce a process model that links the decision-making dynamics involved in adopting algorithmic technologies with two pathways for knowledge work transformation - task alteration and process re-engineering - and the power dynamics implemented to adjust the knowledge regime. Furthermore, we offer propositions that connect these pathways of digital transformation with the outcomes of success and failure in incorporating algorithmic technologies into knowledge work.

In the third essay, "*Enacting MLOps: Definition and underlying theory*", we develop a theoretical model that explains how development and operations practices shape the impact of algorithmic technologies in knowledge work. To achieve this, we conduct an integrative review of how organisations build algorithmic technologies and analyse the resulting outcomes, with a particular emphasis on MLOps, a practice that has gained traction among non-digital native organisations seeking to develop their own machine learning solutions. Drawing on the literature on sociomateriality, we synthesize our findings from the integrative review and suggest ways to develop ethical and effective algorithmic technologies.

By completing these three essays, this thesis seeks to further our understanding of the application of algorithmic technologies in knowledge work, highlighting challenges

and strategic considerations. It also aims to deepen our understanding of how established process improvement methodologies need to be adapted to the context of knowledge work.

Chapter 1

Introduction

The emergence of artificial intelligence and other novel algorithmic technologies has increased the predictive power across many business applications, prompting firms to reorganise work (Frey and Osborne, 2017; Sako et al., 2022) and operations (Mišić and Perakis, 2020) between humans and machines to boost performance. As computer-programmed procedures that emulate human decision-making (Kellogg et al., 2020), algorithmic technologies provide companies with the means to automate knowledge work tasks, such as '*assessing a company's financial and operational risks*', '*designing acidizing treatments for deep wells*' or '*identifying inefficiencies in inventory management*'. In essence, these technologies allow firms to automate work that was previously entrusted solely to expert knowledge workers - those who rely on their mastery of information and expertise to carry out their work (Bechky, 2006; Davenport, 2015; Starbuck, 1992). While scholars have argued that automating such tasks poses potential cost-saving benefits (c.f. Brynjolfsson and McAfee, 2017; Smets et al., 2017), empirical evidence on what happens when companies infuse algorithms into knowledge work remains scarce (Brynjolfsson et al., 2019). Drawing on operations management scholarship, this thesis seeks to redress this issue by studying how companies use algorithms to construct improvement in various contexts of knowledge work and the consequences of those efforts. To that end, the rest of this thesis is structured as follows.

In Chapter 2, we present an empirical background in algorithmic technologies as well as lay the necessary groundwork for this doctoral thesis. This chapter introduces a comprehensive and integrative literature review on the three main streams of literature that this thesis draws upon: technology and occupations, technology, knowledge, and capability, and process improvement and algorithmic technologies. The central focus of this chapter is to understand the mechanisms that may connect the adoption and use of algorithms to consequences on performance in knowledge work. In the latter part of the chapter, we leverage the insights from these bodies of work to underscore the significance of the papers presented in this thesis. Specifically, we introduce the three research themes which are: 1) Can algorithmic technologies improve knowledge work? 2) How do organisations interpret algorithmic technologies to improve knowledge work? and 3) how do development and operations practices affect algorithmic technologies? In doing so, we establish the foundation and motivations of the subsequent chapters.

In Chapter 3, we begin by exploring whether algorithmic technologies can improve knowledge work. To that end, we use a field experiment to analyse the effect of substituting human decisions with machine learning algorithms. The context is inventory management in a large grocery and merchandise retailer. Using a randomized experiment, we study the effects of algorithmic decision-making on planning and control operations. We began by proposing and validating a novel deep learning-based decision model to address the classic problem of inventory record inaccuracy. We then measure the effect of this decision model on key performance indicators, by using an experiment involving 450 large stores in the United Kingdom, and by comparing the performance of stores receiving operations analysts' recommendations and those receiving deep learning recommendations. Our results confirm the benefits of adopting algorithms to address operational problems and provide an important qualification to the moderating

mechanisms at play. Namely, deep learning algorithms improve performance by providing more accurate representations of problems that otherwise, may be ambiguous, for humans. We also show, however, the limitations that need to be considered when deploying new algorithms.

In Chapter 4, we then go on to study the interplay between algorithmic technologies and knowledge improvement across various occupations. This chapter presents a study titled "Using algorithmic technologies in knowledge work" and focuses on unpacking the digital transformation of knowledge work driven by machine learning and other algorithmic technologies. This study is motivated on that while organisations have increasingly adopted algorithms to transform knowledge work, the utility of algorithms as a driver of productivity across diverse types of work remains a point of contention. To address this tension, we examine how organisations interpret and use algorithmic technologies in knowledge work to improve existing processes. The research uses a multiple-case study design and reports on the observation and analysis of the algorithmic adoption in four knowledge work business areas within a multinational energy firm. Cases are built using data using notes from participant observations and from 55 interviews with professionals and subject-matter experts who participated in digital transformation projects involving algorithmic technologies.

This study makes three contributions to the discussion on algorithmic technologies in knowledge-based work. First, the research shows that knowledge workers engage in a dynamic process that allows them to adapt to technological change, increase efficiency in knowledge work processes, and maintain control over the quality of constructed knowledge and learning. Second, the study provides a differentiated view of the existing work on the automation of knowledge work by unpacking variations of work and technology implementation. Lastly, the study proposes a conceptual model that depicts

how organisations use algorithmic technologies to construct improvement opportunities. These findings contribute to a better understanding of the impact of algorithmic technologies on knowledge work.

In Chapter 5, we present a study entitled "Enacting MLOps: How Development and Operations Practices Affect Algorithmic Technologies in Organisations". Departing from the prior chapters, which have investigated how algorithmic technologies affect knowledge work practices, this research inquiry takes a reverse approach and examines how established practices can mould algorithmic technologies. Our focus is on development and operations practices, with particular attention to MLOps, an emergent practice characterized by principles, activities, and artefacts for constructing machine learning algorithms. To explicate the interplay between these practices and algorithmic technologies, the study adopts a sociomaterial perspective that highlights the significance of relationships, actions, and discourse, as well as material objects, in shaping organisational phenomena.

Particularly, this study highlights the importance of understanding the relationships between managers, developers, and users in forming machine learning systems, the performativity of development and operations practices activities and the agentic role of data, automation scripts and other artefacts in shaping normative and epistemic consequences of algorithmic technologies. The study argues development and operations culture can positively influence an organisation's ability to address epistemic concerns, but not normative concerns. It further highlights that addressing such normative concerns requires routinizing activities to check for biases and inequalities. Finally, it proposes that the automation artefacts embedded in development and operations practices may pose short-term challenges but are likely to reduce both

epistemic and normative concerns in the long run. The study offers actionable guidance for scholars and practitioners grappling with the use of algorithms in knowledge work.

In Chapter 6 we present the conclusions of this thesis. Overall, our findings of this thesis seek to contribute to the operations management literature by expanding our theoretical understanding of how algorithms support improvement in the context of knowledge work. These findings can guide firms wishing to automate 'white-collar' work by highlighting how and where algorithmic technology can most effectively support process improvement in this context.

References

Bechky, B.A. (2006). Talking about machines, thick description, and knowledge work. *Organization Studies, 27* (12), 1757-1768.

Brynjolfsson, E. & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review, 98* (4), 1-20.

Brynjolfsson, E., Rock, D. & Syverson, C. (2019). 'Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics', in: Agrawal, A. (Ed), *The Economics of Artificial Intelligence*. University of Chicago Press, Chicago, pp. 23-60.

Davenport, T.H. (2015). 'Process management for knowledge work', in: vom Brocke, J., Rosemann, M. (Eds), *Handbook on Business Process Management*. Springer, Berlin, pp. 17-35.

Frey, C.B. & Osborne, M.A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114* (2017), 254-280.

Kellogg, K.C., Valentine, M.A. & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals, 14* (1), 366-410.

Mišić, V.V. & Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management, 22* (1), 158-169.

Sako, M., Qian, M. & Attolini, J. (2022). Future of professional work: evidence from legal jobs in Britain and the United States. *Journal of Professions and Organization, 9* (2), 143-169.

Smets, M., Morris, T., von Nordenflycht, A. & Brock, D.M. (2017). 25 years since 'P2': Taking stock and charting the future of professional firms. *Journal of Professions and Organization, 4* (2), 91-111.

Starbuck, W.H. (1992). Learning by knowledge-intensive firms. *Journal of Management Studies, 29* (6), 713-740.

Chapter 2

Background and theoretical foundations

2.1 Introduction

The last decade has witnessed an extraordinary interest in algorithmic technologies. Defined as programs that transform inputs into outputs based on a sequence of calculations (Gillespie, 2014), algorithmic technologies are becoming ubiquitous in both day-to-day and organisational life, from matchmaking to hiring job candidates (van den Broek et al., 2021) to selecting investment portfolios (Heaton et al., 2017). Algorithmic technologies have become a source of both excitement and concern due to their ability to imitate human decision-making processes (Brynjolfsson et al., 2019; Krakowski et al., 2022; Lou and Wu, 2021; Raisch and Krakowski, 2021). Such an ability has fuelled the adoption of technology in knowledge work contexts, affecting those who think for a living (Bechky, 2006) including engineers, lawyers (Sako et al., 2022) and physicians (Lebovitz et al., 2022) – occupations traditionally shielded from automation technologies. Adoptions of algorithmic technologies, either symbolic (Pachidi et al., 2021) or substantive (Orlikowski and Scott, 2015a), pose severe consequences for knowledge work, including for the development of expertise (Anthony, 2021), the creation of value and the future of work (Acemoglu and Restrepo, 2019, 2020; Frey and Osborne, 2017; Willcocks, 2020). As the number of knowledge workers continues to increase, with over a billion worldwide scholars are increasingly recognizing the need to investigate how algorithmic technologies are reshaping this critical area of work.

Despite a few notable exceptions (e.g., Ferreira et al., 2016; Sampson, 2021; Senoner et al., 2022; Sodhi et al., 2022; Vanderschueren et al., 2023), research on algorithmic technologies and knowledge work remains nascent in Operations Management, with most studies in the field remaining in manufacturing contexts and focusing on technologies that replace physical labour. This oversight is ironic, given Operations Management has been grounded in practical work since its inception (e.g. Gilbreth, 1911, 1912; Taylor,

1911). Yet, in recent studies, modern work practices and contemporary technologies have been relegated to a secondary, and possibly mundane, role. We argue that this detachment from a thorough understanding of organisational operations, technologies, and work demands redress. Our intention is not to undermine the contribution of traditional research on Operations Management. Rather, our objective is to broaden our understanding of Operations Management by expanding its scope beyond manufacturing technologies and work. This chapter, thus, advocates for a shift in the Operations Management discourse towards an exploration of the role of algorithmic technologies in knowledge work, specifically by examining organisations' day-to-day technologies, processes, and practices.

To advance this debate in Operations Management, we integrate research from other management fields, specifically studies on technology and occupations, technology, knowledge, and capabilities, with work on process improvement. We aim to examine the potential mechanisms that link the adoption and use of algorithmic technologies to knowledge work and operational performance. We found that occupational characteristics, such as shared mental models and levels of authority and discretion among occupational members, can influence the extent to which technology is adopted (Bechky, 2020; Orlikowski, 1992), which in turn affects work practices and organisational structures (Barley, 1986, 1990). Our review also points to that new organisational orders change organisational knowledge capabilities which ultimately affect performance outcomes (Felin and Hesterly, 2007; Grant, 1996; Kogut and Zander, 1992). Finally, we observed that organisations respond to changes in performance by adopting new organisational practices that in turn modify their capabilities (Cusumano et al., 2021; Holweg et al., 2018a; Womack et al., 1990). Together, these findings are integrated into

a theoretical model (see Figure 2-1) that can be used to illuminate our understanding of how algorithms affect improvement in knowledge work.

This paper proceeds as follows. First, we briefly review the concepts of knowledge work and algorithmic technologies, delving into their historical context and definitions to establish the scope of our investigation. We then review the three key bodies of work that serve as foundational sources for this thesis. Extracting and integrating valuable insights from each, we next present three distinctive pathways that warrant further investigation. These pathways serve as focal points for more detailed examinations in subsequent chapters. We conclude by emphasizing the importance of considering the role of algorithmic technologies in knowledge work and the significance of studying modern organisations' day-to-day technologies, processes, and work practices. Recognizing these factors will be pivotal for driving future progress in the Operations Management field.

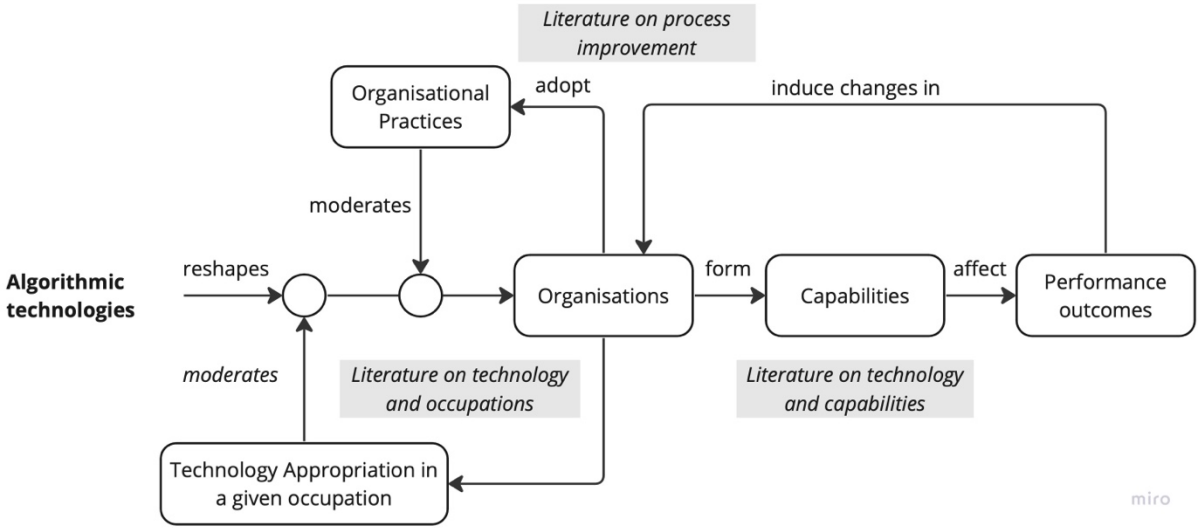


Figure 2-1 Conceptual model of algorithmic technologies on performance outcomes.

2.2 The concept of knowledge work

Over the past half-century, there has been a remarkable transformation from a manufacturing-based economy to one that revolves around knowledge. In this new paradigm, specialized knowledge plays a central role in creating value and securing competitive advantages for organisations (Adler, 2001; Barley and Kunda, 2001; Galbraith, 1967; Kogut and Zander, 1992). This shift towards rationality and scientific thinking has significantly impacted the nature of work, occupations, and markets (Barley and Kunda, 2001). As a result, concepts such as knowledge work (Knights et al., 1993), knowledge workers (Drucker, 1991; Drucker, 2007), and knowledge-intensive firms (Alvesson, 1993; Blackler, 1995; Starbuck, 1992) have become crucial for understanding modern societies.

However, despite their growing importance, these concepts remain shrouded in ambiguity, particularly when it comes to defining knowledge work itself. This ambiguity often creates an aura of mysticism and esotericism around knowledge work, which practitioners leverage to establish their legitimacy and authority in their respective fields (Alvesson, 1993). Even within academic circles, the lack of a precise definition persists, as previous studies frequently introduce the concept of knowledge work without providing a clear definition or describing it solely in terms of what knowledge workers do. Some suggest that knowledge work involves the act of "thinking for a living" (Anthony, 2018; Davenport, 2005) or participating in the knowledge economy (Bechky, 2006). The problem with these definitions is that they either become circular references or make overly broad generalizations, complicating distinguishing knowledge work from other types of work—after all, thinking is inherent in almost all types of work. To address this conceptual ambiguity, we revise the definition of knowledge work delving into its foundational constructs: work and knowledge.

2.2.1 Work

We first started with understanding how management scholars and practitioners employ the term work. Based on our reading, we found that work has often been characterized in various ways. These characterizations include viewing work as encompassing employment relations (Kalleberg, 2009; Spreitzer et al., 2017), serving as the source of workers' identity (Pratt et al., 2006; Rosso et al., 2010), or as the exertion of effort to accomplish a series of tasks (Taylor, 1911). Moreover, in public discourse, work is frequently employed interchangeably with terms such as occupations, jobs, or labour. These uses highlight a core question about work: is work what people say they do for a living (as stated in job descriptions) or the actual practices they engage in?

Organisation theorists have shed light on this question and argue that the nature of work is intertwined with the evolution of technology and organisations (Bailey and Barley, 2020; Barley, 1996; Sergeeva et al., 2020). Consequently, an understanding of work needs to be grounded in specific temporal and geographical contexts, recognising that people with the same job title but operating in different organisations may work in distinctive ways (Barley, 1986). Thus, drawing on the organisation studies, we define inductively work as the concrete, dynamic, and situated activities that individuals do for a living (Barley, 1996; Barley and Kunda, 2001; Orr, 1996).

In this context, we acknowledge that work encompasses a dual nature. On one hand, it involves a series of enacted tasks, facilitated by the utilization of material technologies, aimed at generating economic output (Gilbreth, 1911; Taylor, 1911). On the other hand, it encompasses ongoing accomplishments that shape individual expertise, foster learning, contribute to identity formation, and foster a sense of community (Orr, 1996). Our definition emphasizes that while certain patterns of action may exist across similar job types, these patterns change based on the specific structures, coordination

mechanisms, and technologies that facilitate work. For instance, the nature of work differs for surgeons and nurses depending on whether they operate in a hospital that utilizes endoscopic surgical robots (Beane, 2019; Beane and Orlikowski, 2015; Sergeeva et al., 2020). Additionally, we underscore the significant role that the repetition of activities that encompass work plays in individual learning, work coordination, and the attainment of status within organisations.

2.2.2 Knowledge

The concept of knowledge, much like work, has been subject to diverse interpretations. For instance, Nonaka and Takeuchi (1995, p.58) characterizes knowledge as *'a flow of information, anchored in beliefs and commitment of its holder'*, while Davenport and Prusak (1998, p.5) define it as '*a flux mix of framed experiences, values and contextual information, and expert insight that provide a framework for evaluating and incorporating new experiences and information'*. While these definitions, along with others, often converge on common themes, emphasizing the link between knowledge and action (Choo, 2005; Suchman, 1987; Wiig, 1997), and highlighting the significance of personal interpretation, they also expose substantial disparities that complicate our understanding of knowledge.

These definitions face notable challenges. First, they frequently fall short of elucidating the relationship between context, information, knowledge, and expertise (Tsoukas, 2005). Second, the debate remains regarding whether knowledge is primarily a property of individuals (Polanyi, 1962) or collectives (Wittgenstein, 1953). Lastly, questions persist surrounding the conversion of tacit knowledge into explicit knowledge (Hadjimichael and Tsoukas, 2019; Lebovitz et al., 2021; Nonaka and Takeuchi, 1995). Together, these challenges highlight that an understanding of knowledge should encompass its contextual dimensions, address the ongoing debate concerning individual

versus collective knowledge, and explore the connection between tacit into explicit knowledge.

Considering these challenges, we adopt Tsoukas' (2005) refined definition of knowledge as 'the individual capability to draw distinctions, within a collective domain of action, based on an appreciation of context or theory or both' (p.123). The capacity to make and refine distinctions (Reyes and Zarama, 1998; Vickers, 1965) emphasizes that all knowledge is inherently personal (Polanyi, 1962), as individuals rely on their perceptual and cognitive abilities to make sense of the situation at hand and make distinctions. Additionally, it underscores that knowledge is contingent upon a collective process (Wittgenstein, 1953), as individuals must adhere to shared understandings of standards and appropriateness to draw meaningful distinctions (MacIntyre, 1981; Tsoukas, 2005; Wenger, 1999). This also highlights the crucial role of socialization in the understanding of context (Berger and Luckmann, 1966; Garfinkel, 1984). Furthermore, this perspective emphasizes that tacit and explicit knowledge are intertwined, as the act of contemplating something (focal awareness) necessitates an appreciation of that something within a context or theoretical framework (subsidiary awareness) (Polanyi, 1967; Tsoukas, 2012). In other words, explicit knowledge, which can be articulated and codified, cannot exist without an underlying foundation of tacit knowledge, which is deeply ingrained in an individual's experiences, skills, and intuition.

By embracing Tsoukas' definition, we recognize that knowledge is an individual's capacity to differentiate and discern within a collective sphere of action, relying on an understanding of context and theory. This understanding is shaped by personal perceptual and cognitive abilities, while also being influenced by shared standards and socialization processes. Moreover, it highlights the inseparable nature of tacit and explicit

knowledge, with both contributing to our comprehension and engagement with the world around us.

2.2.3 Knowledge work

Based on the individual characterizations of the concepts of work and knowledge, we propose a combinatorial definition of knowledge work as follows:

Knowledge work refers to the concrete, dynamic, and situated activities that individuals do for a living, characterized by drawing distinctions, within a collective domain of action, based on an appreciation of context or theory or both.

This definition places a central focus on activities as the fundamental unit of analysis, recognizing that while knowledge mastery is inherent to all occupational work to varying extents, certain professions such as lawyers, physicians, and engineers are more prone to engaging in knowledge work compared to occupations like farming, carpentry, and clerical roles. Illustrative instances of knowledge work encompass tasks like diagnosing and treating a patient's illness, evaluating the guilt or innocence of a defendant, or devising an energy recovery system for a new Formula 1 car, whereas non-knowledge work examples entail activities such as precisely cutting wood, planting rice seeds, or transcribing meeting minutes. minutes. The key distinction between knowledge work and non-knowledge work lies in that the former involves the application of discerning distinctions while the latter does not. Importantly, the emphasis on task situatedness highlights that activities that may initially appear as non-knowledge work can still qualify as knowledge work, such as when planting rice seeds involves decisions regarding optimal seed types based on specific soil and seed properties.

Given our definition of knowledge work, we caution against the use of the term '*knowledge workers.*' The increasing trend over the past decades to rationalise work and incorporate production technologies into work, as described by what Barley and Kunda (2001) as the '*technization'* of work, has resulted in occupations encompassing both knowledge and non-knowledge work. As a result, to some extent, everyone can be considered a knowledge worker. Labelling certain occupations as exclusively composed of knowledge workers can have detrimental effects, as it may perpetuate rigid perceptions of work and hinder our understanding of organisations and markets. After all, even professionals like lawyers engage in administrative and mindless tasks alongside their knowledge work.

Accordingly, we propose a revision to Davenport's definition of knowledge-intensive organisations, shifting the emphasis from a high proportion of knowledge workers to a high proportion of knowledge work within the organisation. This redefined concept underscores the significance of the actual activities being carried out rather than solely focusing on the individuals performing them. Moreover, our revised definition introduces a temporal dimension to the concept of knowledge-intensive firms. For instance, a start-up may currently be considered knowledge-intensive due to the prevalence of complex and specialized tasks that demand high expertise. However, as the organisation matures and its processes become standardized, the need for individuals with extensive knowledge may diminish, leading it to become less knowledge-intensive in the future. By making these adjustments, we aim to capture the dynamic nature of knowledge-intensive firms, highlighting the critical role of knowledge work and its potential evolution over time. This updated perspective allows for a more nuanced understanding of organisations' knowledge-intensive nature, considering the fluidity of their requirements and expertise demands as they progress and adapt in their respective industries.

In this thesis, we delve into the interconnected realms of knowledge work, knowledge workers, and knowledge-intensive organisations. However, our focus lies in analysing the intricacies of knowledge work itself. Now, we turn our attention to exploring the second important key concept of this dissertation: algorithmic technologies.

2.3 The concept of algorithmic technologies

2.3.1 A brief history of algorithmic technologies

Searching for machines that can "think"

Contrary to popular perception, algorithmic technologies have a rich and complex history, punctuated by cycles of hype and disappointment (Wooldridge, 2020). The 1956 Dartmouth Conference stands out as a pivotal moment in the history of algorithmic technologies, with the emergence of the field of artificial intelligence (AI), uniting prominent scholars in a quest to understand the potential of '*thinking machines.*' Among these scholars was John McCarthy, whose definition of AI as '(…) *the science and engineering of making intelligent machines'* became a foundation for decades of research in the field (Stone et al., 2022). The conference captured public interest due in part to the presentation of the *Logic Theorist* by Newell and Simon (1956), which represented an algorithm capable of proving mathematical theorems in symbolic logic. However, despite the conference's success in raising public interest in algorithms that imitate human intelligence, it left critical philosophical questions unresolved, including debates around standard methods (Anyoha, 2017), the essence of intelligence and the feasibility of creating machines that can genuinely think (Stone et al., 2022).

During the 1960s, the search for machines that could '*think'* like humans encountered significant obstacles as scholars recognised that real-world problems suffered

from problems of combinatorial explosions, rendering them impossible to solve with the contemporary computing infrastructure. In 1964, Joseph Weizenbaum introduced ELIZA, an early chatbot that utilized natural language technology to simulate the behaviour of a psychotherapist. ELIZA, developed at the MIT Artificial Intelligence Laboratory, attempted the Turing test and highlighted the superficiality of communication between humans and machines (Epstein and Klinkenberg, 2001). While ELIZA's development initially fuelled enthusiasm regarding the potential of algorithmic-assisted patient treatment, the 1970s witnessed the onset of algorithmic technologies, a period known as the *'AI winter.'* In 1973, the British Science Research Council commissioned a seminal report that issued a bleak prognosis for the field, leading many to doubt the legitimacy of Artificial Intelligence algorithms. As Lighthill et al. (1973) observed, *'in no part of the field have the discoveries made so far produced the major impact that was then promised [at the Dartmouth AI Conference].'*

From machines that think to algorithms that learn

In the 1980s, algorithmic technologies prospered in two directions. On one front, scholars made significant strides in developing logic-based algorithms – algorithms based on knowledge or logic, such as expert systems (Feigenbaum, 1981). These algorithms were designed to mimic the decision-making processes of experts by embedding their knowledge into the system, thus creating knowledgeable programmes instead of intelligent ones. While expert systems showed remarkable performance in handling known unknowns, they became brittle and incapable when confronted with scenarios beyond the scope of their programming. The challenges of encoding common sense knowledge into machines are exemplified by the Cyc project (Lenat et al., 1990). The aim was to create a robust ontology that the machine could use to make inferences and understand the world. However, capturing the complexity of the world in an ontology

proved to be a difficult task, as demonstrated by the system's failure to interpret the entity '*FredWhileShaving'* correctly. Instead of understanding that Fred was holding a shaving machine, the system mistakenly inferred that Fred was a machine made of electrical parts (Linde, 1992).

On the other front, algorithmic technologies progressed in systems that emulated learning mechanisms with inference-based algorithms, such as artificial neural networks. These networks enabled computers to learn complex systems by breaking them down into simpler concepts (Goodfellow et al., 2016), following the literature on cybernetics (Hebb, 2005; McCulloch and Pitts, 1943). Researchers developed primitive artificial neural networks with backpropagation algorithms and two or more hidden layers (Rumelhart et al., 1986), laying the groundwork for "deep learning" (Bengio et al., 2006; Bengio and LeCun, 2007; Hinton et al., 2006). Finally, the 1980s also witnessed the launch of Japan's fifth-generation computer system project (1982-1990), a large-scale initiative of the Ministry of International Trade and Industry (MITI) aimed at accelerating the development of massively parallel computers. However, this project ultimately failed to achieve its commercial objectives.

The shift from seeking machines that can think to algorithms that can learn represents a productive ontological change in the study of algorithmic technologies. This change responds to concerns across multiple fields regarding the philosophical and rather subjective nature of thinking and intelligence (Simon, 1995). The philosophical nature of these concepts is captured in questions such as what constitutes genuine intelligence and whether machines should be held responsible and accountable for their actions. The subjectivity inherent in determining whether a non-human entity is thinking is encapsulated in the '*odd paradox,*' which suggests that people often dismiss machine displays that could be perceived as intelligent (McCorduck, 1979). For example, for

many centuries the ability to play chess represented the pinnacle of intelligence, but when IBM's Deep Blue defeated the then-world chess champion, Gary Kasparov, in 1997 (Campbell et al., 2002), Deep Blue was not recognized as an intelligent machine, but rather as one that adhered to a sophisticated set of rules. Thus, a change from thinking machines to learning algorithms has achieved two purposes. First, it focuses algorithms to imitating human decisions as a means to eventually imitate intelligence, recognising that intelligence is an immensely challenging task, requiring the matching of complex workings of the human mind, including factors like speed, autonomy, and generality, all of which are addressed in prior research (Simon, 1995; Stone et al., 2022). Second, in emphasizing learning rather than thinking, the field avoids dismissing the accomplishments of algorithms. Finally, it is noteworthy to mention that while technically the field recognises learning algorithms, public discourse subsumes all algorithms under the banner of artificial intelligence. The adage '*artificial intelligence is done in PowerPoint and machine learning in Python'* comes to mind.

A focus on learning algorithms has demonstrated remarkable progress in imitating human decision-making, with algorithms now exhibiting exceptional performance in specific tasks for which they were trained. For instance, some algorithms can learn from historical data patterns to emulate radiologists' decision-making in detecting breast cancer from mammograms (Hosny et al., 2018). Certain recent algorithms are generative, such as large language models like Chat-GPT and Bard, which can provide answers and create content in a seemingly realistic manner. As such, learning algorithms are increasingly being incorporated into organisational decision-making activities and knowledge work tasks. It appears that we are rapidly approaching a reality where organisational value is created and sustained through the collaborative efforts of human and technological agencies. Moreover, as these learning algorithms become more

autonomous, it is essential to recognize the need to move away from an anthropocentric view of organisations and embrace a post-humanist perspective that acknowledges algorithms as agential counterparts (Anthony et al., 2023). This rise of algorithmic technologies represents the second motivation for this investigation.

2.3.2 Defining algorithmic technologies

In this section, we focus on defining algorithmic technologies by distilling their key features. To begin with, we build upon Gillespie's (2014) definition of algorithms "as encoded procedures for transforming input data into a desired output, based on specified calculations" (p.167) and distinguish between logic-based algorithms and inference-based algorithms. To derive our definitional features of algorithmic technologies we consider the goals that algorithms seek to achieve, as suggested by scholars in the field (Simon, 1995; Russell and Norvig, 2016). We note that when these technologies are used within organisations, they are often intended to mimic the decision-making processes of human experts. Thus, we identify two defining features of algorithms: their ability to i*mitate human decision-making* and the characteristics of the *outcomes* they produce.

The first property that emerges in writings about algorithmic technologies is their ability to imitate human decision-making (Gomez-Uribe and Hunt, 2015; Kellogg et al., 2020), which can be characterised by three dimensions: approach, knowledge type and applicability. This dimension is highlighted, for example, in writings that describe applications of algorithms to automatically select investment portfolios (Heaton et al., 2017) and to automatically identify, characterise and monitor diseases (Hosny et al., 2018). The '*approach'* refers to the methods used by algorithms to simulate human judgement. For instance, logic-based algorithms *encode expert knowledge* into a set of rules and procedures (Feigenbaum, 1981). While inference-based algorithms, such as machine learning models, identify *correlations* in large datasets to make predictions

(Goodfellow et al., 2016; Hollebeek et al., 2016). Now, the "knowledge type" refers to the kind of information utilized by algorithms to imitate decision-making. Logic-based algorithms often rely on incomplete know-how, as experts may find it challenging to fully articulate how they enact their work into explicit knowledge. On the other hand, inference-based algorithms do not require experts to explicitly articulate know-how since they can learn from examples of past events and labelled data, relying on '*know-what*' knowledge (Lebovitz et al., 2021). Finally, '*applicability*' refers to the extent to which algorithmic decision-making can be used in different domains. For instance, Logic-based algorithms are usually suitable for routine tasks but may struggle with novel or unpredictable scenarios, while inference-based algorithms can also be applied to non-routine tasks.

The second property that emerges in writings about algorithmic technologies concerns the characteristics of the outcomes of algorithmic decision-making, including quality, explainability and adaptability. "Quality" concerns whether the standards of algorithmic decision-making outcomes are maintained over time. For example, in inference-based algorithms, the prediction quality changes as data does – a phenomenon sometimes called drift (Bayram et al., 2022; Ditzler et al., 2015; Hofer and Krempl, 2013). "Explainability" concerns whether it is possible to scrutinize and explain outcomes of algorithmic-based decisions. The literature has highlighted that inference-based algorithms are opaque, inscrutable and black boxes (Curchod et al., 2020; Faraj et al., 2018; Hafermalz and Huysman, 2022; Scott and Orlikowski, 2012). Finally, "adaptability" refers to whether the algorithmic outcomes adapt over time to new conditions. Writings on applications of inference-based algorithms, for example, highlight that these algorithms can adjust to new customer preferences (Gomez-Uribe and Hunt, 2015), and discover new patterns on data (Agarwal and Dhar, 2014; Agrawal et al., 2018). A

comparison of these defining features across logic-based and inference-based algorithms is presented below in Table 2-1

Table 2-1 A typology of algorithmic technologies with their defining features

		Logic-based algorithms	Inference-based algorithms
Decision-making	***Approach***	**Codified knowledge** Developers encode complex routines and if-then rules to allow the algorithm to achieve its objective,	**Correlational patterns** Developers connect standard algorithmic architectures with large data sets, to allow the algorithm to create its own processing rules without being explicitly programmed.
	Knowledge type	**Incomplete know-how** Require capturing experts' knowledge into a knowledge base – repository of the information and rules the system uses to make decisions. This may be incomplete as is difficult to articulate know-how even for experts.	**Know-what** Require capturing a large number of examples (and sometimes labels) of what the dec
	Applicability	**Routine tasks** Require standardised inputs; break when encountering unknown unknowns. Algorithms usually only become cost-effective when the automated task is highly repetitive.	**Routine and non-routine tasks** Can work also for non-routine tasks, although the benefit from syntactical and semantic standardisation. Can be cost-effective even in work that is not highly repetitive.
Outcomes	***Quality***	**Static** Outcomes stay static over time unless developers reprogram the algorithm.	**Dynamic** Outcomes change as algorithms get retrained constantly with new data.
	Explainability	**Translucent** Algorithms are black-boxed, but developers can inspect logic-based algorithms code to understand their procedures and outcomes.	**Opaque** Algorithms are black-boxed, and developers cannot inspect decision rules easily, particularly in the case of deep learning algorithms.
	Adaptability	**Limited** Outcomes stay static over time unless developers reprogram the algorithm	**Adaptive** Outcomes change as algorithms get retrained constantly with new data.
	Examples	Expert systems (Feigenbaum, 1981) and Robotic Process Automation, including attended, unattended and intelligent Process Automation (Lacity et al., 2016)	Machine learning algorithms including supervised learning, unsupervised learning, semi-supervised learning, reinforcement learning, deep learning,

	transformers and autoencoders. These algorithms are present in, for example, TripAdvisor hotel ranking algorithm (Scott and Orlikowski, 2012), and eBay online evaluation systems (Curchod et al., 2020).

Examples of algorithmic technologies

For this study, three algorithmic technologies are particularly relevant, namely machine learning, robotic process automation, and intelligent process automation. Machine learning algorithms are inference-based algorithms (Goodfellow et al., 2016), whereby *'a computer is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E*' (Mitchell, 1997). Machine learning applications are developed by data scientists who connect large datasets to training models to learn data patterns and make statistical predictions. Some, not exhaustive, tasks solved with machine learning algorithms are classification, regression, transcription, machine translation, structured outputs, anomaly detection, synthesis and sampling, imputation of missing values, denoising and density estimation (Goodfellow et al., 2016). The most common ML algorithms are supervised, unsupervised, and reinforcement learning. They are often coupled with deep learning architectures.

Robotic process automation (RPA) comprises a category of logic-based algorithms, popularly referred to as 'robots' although they have no physical form, that automate intricate digital processes, often in the back and middle office (Davenport and Brain, 2018). Gartner describes RPA as 'a non-invasive integration technology used to automate routine, repetitive and predictable tasks through orchestrated UI interactions that emulate human actions' (Miers et al., 2019) Programmers build RPA solutions mapping existing processes, crafting rules, and capturing user interface interactions, or connecting

them to Application Programming Interface (APIs). RPA does not necessitate business process redesign and is typically implemented to automate financial processes and transfer data from multiple sources (e.g. emails, structured databases, chats) to transactional systems of record (e.g. ERP systems) (Davenport and Brain, 2018).

Intelligent process automation (IPA), also known as smart RPA, is an advanced category of algorithmic technologies that integrate capabilities from both RPA and ML to imitate human interaction and make sophisticated decisions. Specifically, IPA commonly employs pre-trained machine learning models for computer vision tasks, such as optical character recognition, or voice recognition, such as voice-to-text transcription. The outputs of the machine learning models are subsequently utilized as inputs for RPA tasks.

2.3.3 Summary

In summary, algorithmic technologies are an enduring theme. Our historical review helps us better articulate what distinguishes algorithmic technologies from other technologies and so, as a relevant focus of investigation. We have sought to do so by highlighting distinct aspects of how algorithms imitate *human decision*-making (e.g., approach, knowledge type and applicability) and of its *outcomes* (e.g., quality, explainability and adaptability); identifying and contrasting different types of algorithmic technologies, based on the processing mechanisms; and underscoring the tension between algorithmic technologies and knowledge work. Taken together, these elements provide empirical motivation for recognizing the importance of algorithmic technologies in knowledge work and emphasize the need to integrate them into the study of Operations Management.

2.4 Theoretical foundations

To set the stage for our review of the role of algorithms in improving knowledge work, we briefly lay out an interdisciplinary, yet interrelated, intellectual history of technology in management research. Because disruptive technology can play an important role in rearranging working practices and activities, which restructures firms and markets, we begin by reviewing the literature on technology and occupations. We then go on to summarise the literature on technology, knowledge and capabilities, as organisations that adopt technology modify their current capabilities, coordination mechanisms, and sources of competitive advantage, ultimately impacting organisational performance. Finally, to make this more concrete, we review the literature on process improvement, a perspective that accentuates the vital role of specific practices in reconfiguring work and resources to generate and maintain performance improvement. By integrating these reviews, we aim to identify the mechanisms that link algorithmic technologies with the improvement of knowledge work.

2.4.1 Technology and occupations

The first body of literature relevant to this thesis is on technology and occupations, which emphasizes a symbiotic relationship between technology and occupations. Occupations are socially constructed categories of work, where members share knowledge, work practices and skills as well as social and cultural norms (Abbott, 1988; Anteby et al., 2016; Bechky, 2003; Nelson et al., 2023; Van Maanen and Barley, 1982). Occupational groups, such as engineers, lawyers, and physicians, exemplify this concept. Technology and occupations depend on each other. Technology enables occupational groups to carry out their work practices, shaping the knowledge and skills valued in a particular occupation (Kaplan et al., 2017; Sennett, 2008). For example, video laparoscope

technology enables surgeons to perform precise and minimally invasive surgical procedures, otherwise requiring making large incisions, influencing both surgeons' practice and knowledge (Zetka, 2003). Concurrently, occupational practices lend meaning to technology, legitimizing its use and potential in the occupation (Hutchby, 2001; Leonardi and Vaast, 2017; Orlikowski and Gash, 1994). For example, engineers imbue their drawings with the ideas of power, giving drawings meanings of control over the production line (Bechky, 2003).

The symbiotic nature of the relationship between technology and occupations underscores that technological change results in the restructuring of occupations by altering previously established knowledge, work practices, skills, and power arrangements (Barley, 1986, 1990). Before the industrial revolutions, work was characterized by craft, which was deeply intertwined with an individual's way of life, its rhythms and demands (Barley and Kunda, 2001, p.82). Craftwork involves the mastery of a wide range of embodied skills, and a proclivity toward dedication, communality and exploration (Kroezen et al., 2021). The coordination mechanisms for craftwork relied on hierarchy and community, emphasized by master apprenticeship systems and merchant guilds (Kieser, 1989). These entities regulated access to the craft and safeguarded their members' interests from external threats, including new technologies (Acemoglu and Robinson, 2012; Kellenbenz, 1974; Mokyr, 1998).

In the 18th century, however, the first industrial revolution triggered a transformation of craftwork, characterized by a rationalisation and specialisation of work as well as the separation of work and life. For example, the steam engine and other manufacturing technologies allowed the mechanisation of some artisan work, triggering a fragmentation of craftwork into tasks. This fragmentation facilitated a "deskilling" process (Braverman, 1974; Form, 1987), whereby artisans transitioned from being all-

rounded to specialising in narrower areas (Sorge, 1991). Moreover, the establishment of factories introduced new forms of coordination, such as bureaucratic control and market-based incentives, which brought about new work attitudes, including detachment, individuality, and planning (Kroezen et al., 2021). The separation between work and means of production also marked new forms of segregation in work groups (Marx, 1933), highlighted in categories such as '*blue-collar*' and '*white-collar*,' '*manual*' and '*mental*,' and "*manufacturing*" and "*knowledge*" work (Barley, 1996), resulting in the hierarchizing of occupations and the concentration of power in elite groups.

In general, the advent of technological change prompts the displacement, emergence, and resurgence of occupations. First, the '*displacement effects'* occur when technology automates tasks previously performed by humans, leading to the replacement of labour (Acemoglu and Restrepo, 2020; Autor and Salomons, 2018). For instance, in the Industrial Revolution, manufacturing technologies displaced highly skilled artisans in spinning and weaving (Braverman, 1974; Mokyr, 1992) Second, technological change can result in the emergence of new occupations by '*reinstating*' labour into remaining tasks untouched by automation and new tasks created to support technology (Acemoglu and Restrepo, 2020). In the Second Industrial Revolution, transportation and telecommunication technologies generated new employment opportunities in occupations such as line workers, engineers, and back-office personnel (Chandler Jr, 1977; Landes, 1969; Mokyr, 1992). Finally, technological change can lead to the resurgence of occupations by augmenting productivity and demand for labour - the so-called "productivity effects" (Acemoglu and Restrepo, 2019). Productivity effects have been observed in various industries, including the manufacturing and service sectors, in which technology translated into higher wages and employment growth for some occupations (Graetz and Michaels, 2018)

The impact of technological change on occupations, however, is not deterministic but rather depends on how occupations appropriate technology. The main argument centres on that technology is socially constructed (Leonardi and Barley, 2010; Pinch and Bijker, 1984), and by extension, its consequences for occupational members. For example, two hospitals adopting identical CT scanner technologies arrange social structures in different ways, resulting in distinct configurations of power and task distribution (Barley, 1986, 1990). This perspective underscores the importance of shared mental frames among occupational members in interpreting, sense-making and appropriating new technologies (Bechky, 2020; Gopal and Prasad, 2000; Orlikowski and Gash, 1994). For instance, the response of firearms examiners, toxicologists, and narcotics analysts to the challenge posed by DNA profiling was unique to each group and was largely determined by the shared values and beliefs of their members (Bechky, 2020). This perspective further points to the role of power and status dynamics in mediating the impact of technology on occupations(Anthony, 2018; Barrett et al., 2012; Pine and Mazmanian, 2017). It suggests that individuals with greater authority are better equipped to incorporate technological advancements into their work processes, while those with lower authority face greater challenges in adapting to new technologies.

In summary, technology and occupations are interdependent; technology enables occupational groups to carry out their work practices, shaping the knowledge and skills valued in a particular occupation, while occupational practices lend meaning to technology, legitimizing its use and potential in the occupation. Historically, technological change has resulted in the replacement of many jobs such as bookkeepers and telephone operators (Bresnahan, 1999), particularly those considered routine. But, it has also resulted in the emergence of new occupations in mental-intensive work and increased demand for educated labour, such as engineering and management (Chandler

Jr, 1977; Charles et al., 2013; Goldin and Katz, 1998; Jaimovich and Siu, 2012; Nelson, 1975). While such analyses have been instrumental in identifying labour market trends, it is essential to recognise that technology changes are not deterministic. The interplay between technology and occupations is shaped by various factors, including mental frames, occupational values, institutional contexts, and power and authority arrangements. Occupations, thus, play an active role in shaping technological change. Ways of working that were once perceived as '*primitive*' and inefficient, such as craft, can re-emerge at the will of occupational members (Kroezen et al., 2021). Similarly, technologies that were once displaced by digital and algorithmic alternatives can also make a comeback in response to occupations' nonconformism to the status quo (Nelson et al., 2023).

2.4.2 Technology, knowledge and organisational capabilities

A second body of literature relevant to this thesis is on technology, knowledge and organisational capabilities. A capability is a '*high-level routine[s] (or collection of routines) that, together with its implementing input flows, confers upon an organisation's management a set of decision options for producing significant outputs of a particular type'* (Winter, 2003, p.991). Chandler's (1977) seminal work, coupled with subsequent evolutionary theories of organisational learning (Winter & Nelson, 1982) and the knowledge-based view of the firm (Grant, 1996; Hedlund, 1994; Kogut and Zander, 1992), provides valuable insights into the interplay between technology, knowledge and organisational capabilities. Chandler proposed that technologies are organisational capabilities that arise from a process of learning by trial and error, feedback, and evaluation. He further argued that these capabilities are context-specific and challenging to transfer from one industry or firm to another. This view highlights that organisations construct technology, by building knowledge and expertise around it, shaping unique

identities and value propositions. For example, Toyota and other Japanese firms based their automotive assembly line technologies on Western models but developed their own practices, methods, and organisational structures, leading to the creation of Lean production systems, which gave them an edge over their Detroit competitors (Fujimoto, 1999; Holweg, 2007; Womack et al., 1990).

A perspective on knowledge, technology, and capabilities recognizes knowledge as the primary determinant of firm performance (Grant, 1996). This perspective posits that knowledge is a social accomplishment that provides organisations with a means to compete (Kogut and Zander, 1992). Knowledge is social in that it is embedded in the relationships among organisational actors, particularly in informal communities with shared interests and practices (Brown and Duguid, 1991). Therefore, a firm's performance hinges on its ability to efficiently and effectively (re)combine individual knowledge into collective knowledge (Grant, 1996; Leonard-Barton, 1998). This body of literature distinguishes between explicit knowledge, which can be easily articulated and shared among members, and tacit knowledge, which is not easily expressed but enacted through organisational actors' expertise, skills, and practices over time (Winter and Nelson, 1982). While formal mechanisms such as patents can protect explicit knowledge from imitation, its easy sharing limits its potential as a source of competitive advantage for organisations (Winter, 2009). In contrast, tacit knowledge, which is not easily learned, is a superior source of competitive advantage (Dierickx and Cool, 1989), as its embeddedness in social structures makes it rich and challenging to imitate or substitute (Barney, 1991).

In summary, extant research on technology, knowledge and capabilities underscores the centrality of knowledge in shaping organisational performance. This literature suggests that knowledge is often challenging to replicate, whether intentionally by the

firm or inadvertently through imitation and unwanted diffusion. Moreover, this perspective accentuates that knowledge "can provide sources of competitive advantage because it depends on the unique interrelationships between people, routines, and technologies that are highly inimitable" (Lei et al., 1996, p.552). Consequently, the dynamic interplay between knowledge and technology is crucial, as it has the potential to impact both performance and inertia and either facilitate or hinder organisational adaptability.

2.4.3 Process Improvement

A third body of literature relevant to this thesis is on process improvement, which has provided fascinating and practical insights into organisations' value-creation and the sources of efficiency. As such, it may help us explore the role of algorithms in improving knowledge work. We first briefly review the concept of process improvement and key approaches (e.g., Lean, Six Sigma, and Agile). Process improvement refers to 'bundles' of complementary organisational practices that aim at creating and sustaining a positive change in the efficiency, effectiveness, or robustness of business processes. The notion of complementary 'bundles' stresses that organisations need to deploy inter-related practices; the application of dislocated tools yields no performance benefits and leads to fatigue (Holweg et al., 2018a; Macduffie, 1995; Osterman, 1994; Radnor et al., 2012). Process improvement stems from organisational change that drives and sustains superior performance, rather than from small temporal and localised benefits. That means that improvement requires substantive changes in the actions and beliefs of individuals who carry out the process to avoid slipping back to the status quo ante, or worse (Collins and Browning, 2019; Holweg et al., 2018b)

Process improvement has been important in operations management research for more than a century (Holweg et al., 2018a), though with rather different trajectories.

Classic works stretch back to the late 19th and early 20th centuries, introducing scientific (Taylor, 1911), administrative (Fayol, 1916), and bureaucratic (Weber, 1930) management concepts; and discussing their effects on process and structural efficiency within organisations. The application of time and motion studies (Gilbreth, 1911, 1912), the 'separation between hand and brain' and experimentation to find 'the best way' (Taylor, 1911) helped managers improve work time appraisal, labour wage-rate setting, and appropriate work-method determination. Then, Shewhart (Shewhart, 1926), Deming (Shewhart and Deming, 1939), Juran, Ishikawa, and Shingo popularised the idea of continuous improvement in the 1980s and 1990s under the rubric of Just-in-Time, Lean, Six Sigma and Agile (Ōno, 1988; Schonberger, 1982, 1987; Womack et al., 1990). During the last three decades, other theories of process improvement have emerged including, the sand cone model (Ferdows and De Meyer, 1990), Swift Even Flow (Schmenner, 2012; Schmenner and Swink, 1998) and Process Theory (Holweg et al., 2018a). For a historical review see 'Improving processes' (Holweg et al., 2018a)

Key process improvement approaches

Over time, scholars and practitioners have advanced diverse process improvement approaches that reproduce 'bundles' of harmonious practices to drive superior performance (Shah and Ward, 2003; White et al., 1999). Notable approaches include Lean, Six Sigma and Agile, with the Theory of Constraints (Goldratt, 1997; Goldratt and Cox, 1984), Business Process Reengineering (Davenport, 1993; Davenport and Short, 1990; Hammer and Champy, 1993) and 'Good to Great' (Collins, 2001) also gaining attention. Each approach highlights the role of specific management practices in driving and sustaining improvement. Theory of constraints highlights practices for identifying and handling processes limiting factors like bottlenecks or 'constraints' (Goldratt, 1997;

Goldratt and Cox, 1984); Business Process Reengineering highlights practices for rethinking and redesigning process workflows to challenge institutionalised assumptions of traditional processes and achieve breakthrough improvements (Hammer and Champy, 1993); and 'Good to Great' highlights practices for cultivating disciplined people, thinking, and action, such as hiring the right talent and keeping them focused on thinking and acting only upon important and urgent activities (Collins, 2001).

Parallel to studying management practice 'bundles', scholars have noted process improvement approaches play a role in how firms execute their strategy: they reconfigure resources and capabilities, repositioning the firm on the market. The firm's ability to compete is improved when managers select initiatives that align process capabilities with market needs (Hayes and Wheelwright, 1979). Such alignment hinges on the degree of standardisation and repetition of the process – what this article refers to as the nature of the process. For example, Lean, Six Sigma, and Agile have their respective main domains of application in the Product/Process Matrix (see Figure 2-2). In the following, we review the three main current approaches for process improvement.

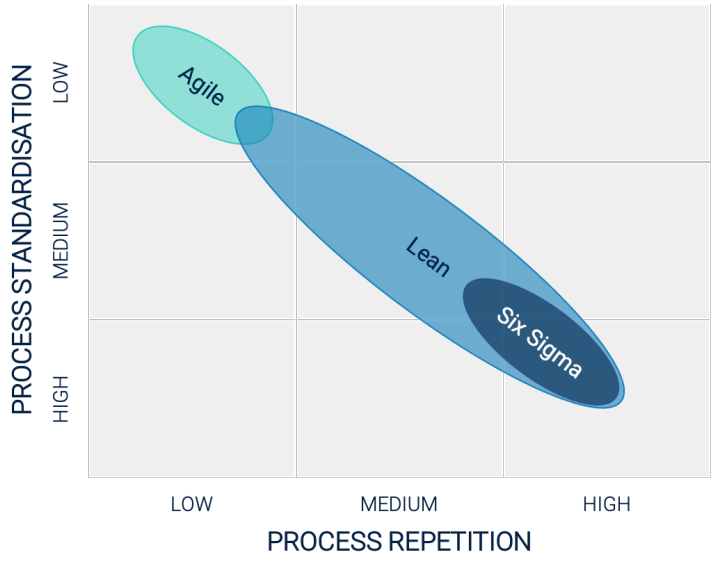


Figure 2-2 Process improvement approaches. The main domains of application in the Standardisation and Repetition Matrix. Source: Based on Six Sigma lecture slides (Holweg)

Lean

Lean is an abstraction for the philosophy, the principles (Hines et al., 2004; Womack and Jones, 1996), and the bundle of practices (Shah and Ward, 2003, 2007) used in the Toyota Production System (Cusumano et al., 2021; Hall, 1983; Schonberger, 1982; Womack et al., 1990). The objective of Lean is to eliminate/reduce the waste (Muda), overburden (Muri) and unevenness (Mura) of a process (Bicheno and Holweg, 2000). Within this body of work, scholars emphasize a democratic approach to process improvement, underscoring the importance of teaching and empowering workers to identify opportunities for improvement in their daily operations (Bicheno and Holweg, 2000). This approach relies on mental models of 'waste' (or *muda* in Japanese), which refers to non-value-adding steps in the production process. Eight wastes are commonly cited in the literature, namely transportation, inventory, motion, waiting, overproduction, over-processing, defects (Ōno, 1988), and non-utilised talent (Womack and Jones, 1996). While it sounds like a mundane activity, removing waste can lead to substantial competitive advantages when this is performed continuously and appropriately. Perhaps the most cited example is how Toyota outperformed automotive manufacturers in the United States by decreasing inventory and removing other wastes (Hall, 1983; Schonberger, 1982; Womack et al., 1990)

Lean is a post-rationalisation of a long and dynamic learning process (Cusumano, 1985; Fujimoto, 1999; Holweg, 2007); Lean *'best practices'* are a recompilation, abstraction, and adaptation of the managerial practices that helped the Japanese automotive manufacturers to cope with a complex post-war environment and transform themselves into dominant players in the industry (Fujimoto, 1999). Thus, although Lean yielded successful productivity gains for Toyota and Nissan, it is not a silver bullet for

process improvement; Lean exhibits path dependency (Holweg, 2007), and its productivity gains and stickiness are mediated by contextual factors (Shah and Ward 2003, Danese, Romano, and Boscari 2017). Yet, Lean yields better results in processes that exhibit a medium degree of standardisation and repetition. This makes it a very versatile approach that works for a range of manufacturing and service processes including software development (Staats et al., 2011), aircraft production (Browning and Heath, 2009) and healthcare services (Radnor et al., 2012).

Six Sigma

Six Sigma is a methodology to improve processes, products, and services that origins from the total quality management (TQM) tradition by leveraging statistical techniques to monitor and control variations in quality and relies on improvement specialists to reduce process variation Its history and main approaches have been discussed by Schroeder et al. (2008) and others (Antony, 2006; Barney, 2002; Kwak and Anbari, 2006). Schroeder et al. (2008) define Six Sigma as '*an organized, parallel-mesostructure to reduce variation in organisational processes by using improvement specialists, a structured method, and performance metrics to achieve strategic objectives'*. These parallel structures *'operate outside of, and do not directly alter, an organisation's normal way of operating'* (Lawler III, 1996), and enable organisations to identify and eliminate defects, minimize variability, and ultimately enhance performance Examples of applications of Six Sigma have been documented in manufacturing firms (Pande et al., 2001), healthcare services (Taner et al., 2007; Thomerson, 2001), and financial services (Antony, 2006; Rucker, 2000).

Six Sigma usually yields improvement results for processes with a high degree of repetition and a high degree of specification because it relies on statistical techniques to

monitor and control its quality. Research indicates a focus on specifying and monitoring customer-oriented and financial performance metrics (Schroeder et al., 2008); an approach in which process improvement is driven by engagement with customers and detailed specification of the critical-to-quality requirements (Evans and Lindsay, 1989). This strong focus on specifying performance metrics is supported by applying a systematic approach to problem-solving based on the PDCA[1], DMAIC[2] and DFSS[3] (Schroeder et al., 2008). Furthermore, Six Sigma provides training to the organisation to create a common language, institutionalises process improvement and creates structural control (Eisenhardt, 1985; Ouchi, 1979; Schroeder et al., 2008). This investment requirement is usually compensated when the repetition and standardisation of a process are high enough to create an attractive return on investment.

Agile

The concept of Agile first appeared in 1991 at the Iacocca Institute (Nagel and Dove, 1991) and refers to *'an organisation that displays the ability to thrive and reconfigure itself'* (Krishnamurthy and Yauch, 2007) *'to provide customer-driven products and services in a fast-changing market environment'* (Yusuf et al., 1999). This paradigm contrasts with mass production and instead focuses on manufacturing flexibility achieved through information technology (IT) and the networked economy (Burgess, 1994). Previous research on Agile has studied its relationship with Lean (Inman et al., 2011). The predominant view is that Agile is the ultimate step for Just-In-Time (JIT) and Lean manufacturing (Inman et al., 2011; Narasimhan et al., 2006). However, other studies

[1] PDCA: Plan, Do, Check, Act
[2] DMAIC: Define, Measure, Analyse, Improve, Control
[3] DFSS: Design for Six Sigma

characterize Agile and Lean as mutually exclusive concepts (Harrison, 1994), while others find them to be mutually supportive (Ramesh and Devadasan, 2007).

Compared to Lean, Agile can be applicable in processes with a low degree of standardisation and repetition. In Ferdows and De Meyer's (1990) parlance, Agile approaches process improvement as a case of building quality and meeting customer specifications. Recent research on this stream shows that additive manufacturing technologies used in Agile production enable mass customisation, increasing the organisation's capability to meet customer requirements (Baumers and Holweg, 2019). Similarly, Agile has been applied in other domains as a means to improve product development, software development (Beck and Fowler, 2001; Poppendieck and Poppendieck, 2003), and entrepreneurship processes (Ries, 2011) - all processes that exhibit a low degree of standardisation, and need to be continuously readjusted.

Process improvement theory: Practices and mechanisms

Up to this point, this research has presented three process improvement approaches that have been pivotal for delivering excellent products or services, and ultimately for firm performance. To make possible the discussion of how process improvement theory can inform the nascent debate of algorithms and knowledge work productivity, this study lays out its assumptions about the relationship between process improvement practice '*bundles'* and heterogeneity. For a long time, management scholars have been concerned with addressing the question of where the primary source of heterogeneity resides to explain differences in productivity between firms, hence the source of competitive advantage. Heterogeneity exists at all levels, such as business, corporate and industry levels (McGahan and Porter, 1997; Rumelt, 1991) - yet operations management and

knowledge-based scholars argue that most heterogeneity resides at the firm level (Adler and Cole, 1993; Bloom et al., 2012, 2016).

Specifically, concerning process improvement theory, scholars agree that firm-level heterogeneity stems from managerial practices that drive improvement in labour productivity, quality, efficiency and costs alike (Shah and Ward, 2003; White et al., 1999). Thus, process improvement approaches work because they reproduce a '*bundle'* or '*system'* of complementary practices that create harmonious effects to drive superior performance. For instance, (Shah and Ward, 2003) explain these practices *'work synergistically to create a streamlined, high-quality system that produces finished products at the pace of customer demand with little or no waste".*

Although most research has been conducted in global manufacturing, evidence shows that differences in management practices also explain heterogeneity beyond production, for example in major projects (Holweg and Maylor, 2018), services (Ahlstrom, 2004; Gupta et al., 2016), and technology adoption (Bloom et al., 2012). Particularly of interest for this research is Bloom and colleagues' (2012) study of US and non-US multinationals operating in the UK which shows that management practices explain better IT technology appropriation and firm performance.

Process improvement theory: Contextual factors

To extend process improvement theory beyond manufacturing contexts and study how 'bundles' of practice shape the impact of algorithms in knowledge, this section briefly discusses criticisms and revisits theoretical underpinnings, assumptions and levels of analysis. For our understanding, the criticism to extend process improvement theory to explain new phenomena is grouped around three arguments. First, most process improvement practices are grounded in manufacturing (c.f. McLachlin, 1997; Shah and

Ward, 2003), rather than in knowledge work contexts. For example, cellular manufacturing and preventative maintenance practices, which are rooted in the JIT[4] and TPM[5] '*bundles*', make little sense for knowledge-based organisations, such as law and accounting firms. Second, research on explaining practices as a source of improvement has mostly focused on efficiency rather than effectiveness logic, which is more important for service provision. For example, one cares more about how many patients a doctor successfully treats, rather than how many patients a doctor sees per hour. Yet, most prominent theories in the field focus on practices for reducing throughput time or undesired variability (Holweg et al., 2018a; Schmenner and Swink, 1998). Third, these practices were conceived for repetitive work, which questions their applicability in firms where work is less specified (Collins and Browning, 2019; Staats et al., 2011). These three points have led some scholars to argue that process improvement theory is incompatible with new phenomena.

This study argues differently. While some scholars see theoretical incompatibility, this research reasons that such conceptualization stems from a locus-related misattribution of the source of improvement. In other words, process improvement practices work because of the principles they embody, rather than because of the practices themselves. As presented earlier, Agile, Lean and Six Sigma are approaches that embody different practices, and nonetheless, they drive productivity improvement alike in their respective domains of applications. These practices, for example, argue for removing unnecessary steps as well as minimising process unevenness. Thus, those who argue for theoretical incompatibility focus on operational practices and miss that process

[4] JIT: Just in Time

[5] TPM: Total Preventive Maintenance

improvement theory is about the principles that precede and explain causal mechanisms of performance driven by those contextualised practices.

Specifically, in lieu of this alternative understanding, there are two implications for extending process improvement theory to explain the interplay between algorithms technologies and knowledge work. First, it requires understanding the contextual characteristics of knowledge work and second, moving the level of analysis from studying operational practice '*bundles*' to organisational ones.

2.5 The interplay between algorithmic technologies and knowledge work

Having reviewed the three distinct bodies of knowledge, this section seeks to integrate our review of algorithmic technologies and knowledge work with the three bodies of literature presented in the previous section. Based on this integration, we identify three research questions that require further investigation and shall be the focus of the forthcoming chapters. Although these questions are not intended to be exhaustive, they illustrate how the field of Operations Management can be enriched by incorporating an understanding of algorithmic technologies and knowledge work into its investigations. Answering these research questions, we argue, can develop a nuanced understanding of the impact of algorithmic technologies on knowledge work, establish the factors that influence the adoption and appropriation of algorithms, and identify the practices that enable knowledge work organisations to improve performance.

2.5.1 Research theme 1: Can algorithmic technologies improve knowledge work?

The first research theme that emerges from our integrative review is to confirm whether algorithms indeed improve knowledge work. Algorithmic technologies possess unique characteristics that make them a promising area for investigating improvements in knowledge work. They provide a means for organisations to emulate decision-making processes (Brynjolfsson and Mitchell, 2017; Frey and Osborne, 2017; Kellogg et al., 2020; McAfee and Brynjolfsson, 2012), while also enabling organisations to encapsulate knowledge into digital artefacts at a low cost. This can yield numerous benefits, including easier distribution of knowledge across different regions, economies of scope, and labour cost savings through automation and arbitrage (Lacity et al., 2016).

Recent investigations on organisational learning, however, have questioned the extent to which algorithms improve knowledge work. For one, in knowledge work contexts, '*the critical input in production and primary source of value is knowledge*' (Grant, 1996). Any alterations to the source of knowledge, then, may have a significant impact on firm performance. Critics of algorithmic technologies caution against the differences between human and algorithm learning mechanisms and their impact on knowledge creation and maintenance. Humans learn through interactions with the environment (Ingram and Baum, 1997; Levitt and March, 1988) as well as through socialization and interactions with others, driven by personal motivations, creative impulses, and perceptions (Maitlis and Christianson, 2014; Weick, 1988, 1993). Contrary, algorithms rely on formal rationality and lack both sentience and contextual knowledge. Thus, some studies argue algorithms may lead to detrimental consequences for the firm, such as exacerbating myopic learning (Balasubramanian et al., 2022; Choudhury et al., 2020). This argument is summarised by Mintzberg's (1994) discussion of technologies and other formal systems. He argues that *'formal systems could certainly process more*

information, at least hard information. But they could never internalize it, comprehend it, synthesize it' (p.19).

This tension between algorithmic technologies and improvement in knowledge work needs further investigation. Some recent studies suggest algorithms create improvement by transforming large and complex datasets into insights otherwise missed by human experts (Bernstein et al., 2019; Ferreira et al., 2016; Morales and Wang, 2010). For example, Senoner et al. (2022) demonstrates that explainable artificial intelligence can facilitate organisations to prioritise improvement activities in semiconductor manufacturing. While these studies have made a valuable contribution, our understanding of the effectiveness of algorithmic technologies is incomplete, as most studies assume that organisational actors will fully adhere to algorithmic suggestions. In practice, technology appropriation is influenced by microdynamics, which may trigger processes of algorithmic aversion (Dietvorst et al., 2015) and hinder organisations from realizing the performance benefits of algorithmic technologies. Therefore, a more nuanced understanding of algorithmic technologies in practice and their consequences for operational performance in knowledge work is necessary.

Chapter 4 explores in detail the impact of algorithmic technologies on knowledge work by presenting the results of a randomized control trial conducted on a large retailer. The experiment aimed to measure the impact of algorithmic recommendations on various key performance indicators by replacing an expert operation analyst with algorithmic recommendations in a specific task. Our study allows us to determine whether knowledge work performance is better served by human expertise, which possesses tacit knowledge that is challenging to encode into algorithms (c.f. Balasubramanian et al., 2022). To minimize the risk of affective or rational aversion, we conducted the experiment without

revealing whether the recommendations originated from a human or an algorithm. The proposed study's conceptual framework is presented below in Figure 2-3.

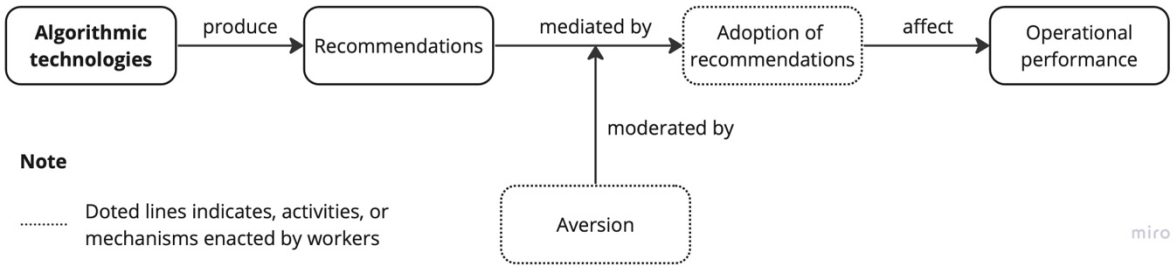


Figure 2-3 Conceptual model for studying research theme 1

2.5.2 Research theme 2: How do organisations use algorithmic technologies to improve knowledge work?

The second research theme that warrants further examination is the transformative processes that enable or inhibit the improvement of knowledge work. Particularly, our understanding remains incomplete regarding how occupational groups interpret and appropriate technology to improve knowledge work. The narratives in the literature on technology and occupations agree that algorithmic technologies primarily affect knowledge work (Sampson, 2021; Susskind, 2020; Susskind and Susskind, 2016) but the debate on the precise effects of these technologies, with some highlighting that 47% of jobs are at high risk of being computerised (Frey and Osborne, 2017) and others calling to cancel such robo-apocalypse narratives (Willcocks, 2020). More sophisticated explanations suggest that shared values (Bechky, 2020), status (Anthony, 2018), and power (Barrett et al., 2012) play a role in the appropriation of algorithmic technologies. A recent examination in the context of legal work has further demonstrated that professional jurisdictional control affects legal work in the US and the UK (Sako et al., 2022). These views highlight that while algorithmic technologies are, and will continue,

driving the restructuration of knowledge work, their precise effects need to be studied considering institutional arrangements, underscoring a micro perspective of technological change.

Similarly, the literature on technology, knowledge and capabilities suggests that analysing the microfoundations of knowledge can be beneficial for our understanding of how algorithmic impact performance. Microfoundations are about 'locating (theoretically and empirically) the proximate causes of a phenomenon (or explanation of an outcome) at a level of a level of analysis lower that of the phenomenon itself' (Felin et al., 2015, p.586). Organisational learning and knowledge-based theorists have embraced this approach to cognise the link between human knowledge and value. Examples of micro levels of analysis used in previous research are individuals (Felin and Hesterly, 2007; Gavetti, 2005), skills, processes, procedures and organisational structures (Teece, 2007).

Per these two literatures, in Chapter 5, we propose an investigation of the transformation of knowledge work considering tasks as our level of analysis. Such an investigation involves paying attention to the day-to-day activities and micro dynamics involved in the appropriation of technology and the construction of knowledge. We do so by studying how the tasks in knowledge work change as organisations embrace or resist algorithmic technologies. Focusing on tasks is consistent with prior research on labour markets, which has used the occupational tasks captured in the Dictionary of Occupational Titles (DOT) or the O*NET databases to study the impact of technology on routine, non-routine, manual, and information processing work (Autor, 2013; Autor et al., 2003). One key difference between these studies on labour markets and ours is that paradoxically labour market studies on tasks analysis remain at the macro level, which has trapped them in generic abstractions of work. Contrary, we suggest in Chapter 5 a micro approach to examining the nature of the task or content of work in practice. A

micro-analysis of tasks is critical because it highlights the embeddedness of work in wider organisational systems, underscores that the enactment of tasks requires knowledge of workers' knowledge to be executed, and highlights that changes in tasks lead to changes in performance. This can help to move the literature beyond work abstractions and focus on analysing questions of variance of work across organisations (Bailey and Barley, 2020). A conceptual model for studying this transformation is presented in Figure 2-4.

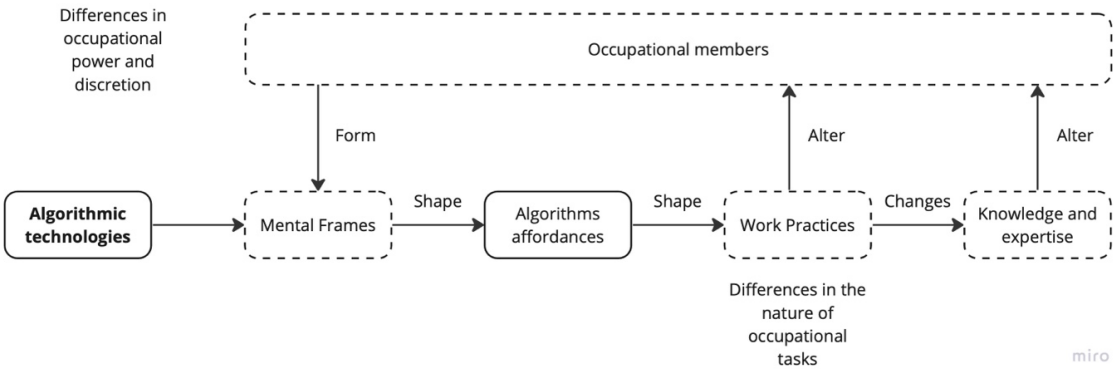


Figure 2-4 Conceptual model for understanding the transformation of knowledge work.

2.5.3 Research theme 3: How do development and operations practices affect algorithmic technologies?

The third research theme that requires further examination concerns the influence of development and operations practices on algorithmic technologies and their implications for organisations. The existing literature on process improvement suggests that managerial practices can be a source of heterogeneity in organisational performance. This perspective is supported by research on technology appropriation, which indicates that management practices play a crucial role in reaping benefits from technology investments (Bloom et al., 2012; Brynjolfsson et al., 2019).

For research on algorithmic technologies in knowledge work, this means that the performance outcomes of algorithms are less deterministic than previously thought and that managerial practices play a significant role in enhancing or inhibiting the performance of algorithmic technologies. However, it is important to note that this research focuses on knowledge work, not manufacturing work. Therefore, it is necessary to consider the contextual differences between these two forms of work, as emphasized by critics of process improvement literature. To address these concerns, we suggest that research should shift away from imposing the Lean and Six Sigma bundle of practices over new phenomena and instead concentrate on analysing the activities performed by those who develop and operate algorithmic technologies, as well as the context in which these practices are carried out. This approach can provide a more grounded and practical understanding of how development and operations practices can enhance algorithmic technologies in knowledge work.

In Chapter 6, we aim to conduct a theoretical exploration of the design and operations of algorithmic technologies. To ensure that our analysis remains practical and relevant to actual organisational practices, we will focus on MLOps, an emerging practice in computer science and information systems to develop machine learning applications that are gaining popularity among non-digital native organisations (Benbya et al., 2020; Vartak, 2022). To bridge the gap between theory and practice, we advocate for the application of the sociomateriality lens (Jones, 2014; Orlikowski and Scott, 2008, 2015b), which emphasizes the role of both material and social factors in shaping daily activities. This lens is particularly relevant for our investigation since it takes a practice-oriented approach. We propose using this lens to examine how the design and operations of algorithmic technologies influence their outcomes in organisations.

2.6 Conclusion

In this chapter, we have laid essential groundwork for studying algorithmic technologies and knowledge work in Operations Management. Our exploration began with an extensive review of interdisciplinary literature, emphasizing the significance and prevalence of knowledge work and algorithmic technologies in modern organisations. One crucial observation was that algorithmic technologies can simulate decision-making activities, a fundamental aspect of knowledge work, suggesting their potential to transform it. However, we also recognized that algorithmic technologies possess specific characteristics related to the quality, explainability, and adaptability of their knowledge outcomes. These characteristics raise important questions about whether their use in transforming knowledge work can genuinely lead to improved operational performance for organisations. To provide the theoretical foundation for these questions, we have reviewed three bodies of work related to technology, occupations, knowledge, capabilities and improvement. This review unveiled three research themes that warrant further exploration in the context of algorithmic technologies and knowledge work.

The first research theme focuses on the extent to which algorithms can enhance knowledge work tasks. While algorithmic technologies have the potential to improve decision-making processes, facilitate knowledge distribution, and reduce labour costs, their formal rationality and lack of contextual knowledge may lead to adverse performance consequences. This tension motivates further research into the effectiveness of algorithmic technologies in improving knowledge work.

The second theme examines how organisations interpret algorithmic technologies to improve knowledge work. While algorithmic technologies drive the restructuring of knowledge work occupations, the precise effects of these technologies remain unclear, depending on occupational characteristics and contextual factors that underpin different

institutional arrangements. Hence, further research is required to examine the connection between algorithmic technologies, the pathways of transformation of knowledge work, and their associated impact on operational performance.

Finally, the third theme explores the development and operations practices that can enable or inhibit the improvement of adopting algorithmic technologies in knowledge work. As technology reshapes organisational practices, it is equally important to understand how these practices, in turn, shape the effects of algorithmic technologies. This theme highlights the need to explore how organisational practices influence the outcomes of algorithmic technologies within an organisation.

In the forthcoming chapters, we will explore these three research themes in-depth, to advance our understanding of Operations Management in the age of algorithmic technologies. By gaining a more grounded and current appreciation of the work and technologies driving the contemporary economy; identifying the mechanisms linking algorithmic technologies to work practices, and performance outcomes; and embracing a theoretical understanding that recognizes the central role of algorithmic technologies and knowledge work, we hope to contribute to the progress of this critical field.

References

Abbott, A.D. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press, Chicago.

Acemoglu, D. & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *The Journal of economic perspectives, 33* (2), 3-30.

Acemoglu, D. & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy, 128* (6), 2188-2244.

Acemoglu, D. & Robinson, J.A. (2012). *Why Nations Fail : The Origins of Power, Prosperity and Poverty*. Profile Books, London.

Adler, P.S. (2001). Market, hierarchy, and trust: The knowledge economy and the future of capitalism. *Organization Science, 12* (2), 215-234.

Adler, P.S. & Cole, R.E. (1993). Designed for learning: A tale of two auto plants. *Sloan Management Review, 4*, 85-89.

Agarwal, R. & Dhar, V. (2014). Big data, data science, and analytics: The opportunity and challenge for IS research. *Information Systems Research, 25* (3), 443-448.

Agrawal, A., Gans, J. & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press, Boston, MA.

Ahlstrom, P. (2004). Lean service operations: translating lean production principles to service operations. *International Journal of Services Technology and Management, 5* (5-6), 545-564.

Alvesson, M. (1993). Organizations as rhetoric: Knowledge-intensive firms and the struggle with ambiguity. *Journal of Management Studies, 30* (6), 997-1015.

Anteby, M., Chan, C.K. & DiBenigno, J. (2016). Three lenses on occupations and professions in organizations: Becoming, doing, and relating. *Academy of Management Annals, 10* (1), 183-244.

Anthony, C. (2018). To question or accept? How status differences influence responses to new epistemic technologies in knowledge work. *Academy of Management Review, 43* (4), 661-679.

Anthony, C. (2021). When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative Science Quarterly, 66* (4), 1173-1212.

Anthony, C., Bechky, B.A. & Fayard, A.-L. (2023). "Collaborating" with AI: Taking a System View to Explore the Future of Work. *Organization Science, Forthcoming*

Antony, J. (2006). Six sigma for service processes. *Business Process Management Journal, 12* (2), 234-248.

Anyoha, R. (2017). *The History of Artificial Intelligence*, Science in the News. Harvard University, The Graduate School of Arts and Sciences

Autor, D. & Salomons, A. (2018). Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brooking papers on economic activity,* (Spring 2018), 1-63.

Autor, D.H. (2013). The "task approach" to labor markets: an overview. *Journal for labour market research, 46* (3), 185-199.

Autor, D.H., Levy, F. & Murnane, R.J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics, 118* (4), 1279-1333.

Bailey, D.E. & Barley, S.R. (2020). Beyond design and use: How scholars should study intelligent technologies. *Information and Organization, 30* (2), 100286.

Balasubramanian, N., Ye, Y. & Xu, M. (2022). Substituting human decision-making with machine learning: Implications for organizational learning. *Academy of Management Review, 47* (3), 448-465.

Barley, S.R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly, 31* (1), 78-108.

Barley, S.R. (1990). The alignment of technology and structure through roles and networks. *Administrative Science Quarterly, 35* (1), 61-103.

Barley, S.R. (1996). Technicians in the workplace: Ethnographic evidence for bringing work into organizational studies. *Administrative Science Quarterly, 41* (3), 404-441.

Barley, S.R. & Kunda, G. (2001). Bringing work back in. *Organization Science, 12* (1), 76-95.

Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management, 17* (1), 99-120.

Barney, M. (2002). *Motorola's second generation*, Six Sigma Forum Magazine, pp. 13-16.

Barrett, M., Oborn, E., Orlikowski, W.J. & Yates, J. (2012). Reconfiguring boundary relations: Robotic innovations in pharmacy work. *Organization Science, 23* (5), 1448-1466.

Baumers, M. & Holweg, M. (2019). On the economics of additive manufacturing: Experimental findings. *Journal of Operations Management, 65* (8), 794-809.

Bayram, F., Ahmed, B.S. & Kassler, A. (2022). From concept drift to model degradation: An overview on performance-aware drift detectors. *Knowledge-Based Systems, 245*, 108632.

Beane, M. (2019). Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly, 64* (1), 87-123.

Beane, M. & Orlikowski, W.J. (2015). What difference does a robot make? The material enactment of distributed coordination. *Organization Science, 26* (6), 1553-1573.

Bechky, B.A. (2003). Sharing meaning across occupational communities: The transformation of understanding on a production floor. *Organization Science, 14* (3), 312-330.

Bechky, B.A. (2006). Talking about machines, thick description, and knowledge work. *Organization Studies, 27* (12), 1757-1768.

Bechky, B.A. (2020). Evaluative spillovers from technological change: The effects of "DNA envy" on occupational practices in forensic science. *Administrative Science Quarterly, 65* (3), 606-643.

Beck, K. & Fowler, M. (2001). *Planning extreme programming*. Addison-Wesley Professional,

Benbya, H., Davenport, T.H. & Pachidi, S. (2020). Special issue editorial: Artificial intelligence in organizations: Current state and future opportunities. *MIS Quarterly Executive, 19* (4), IX-XXI.

Bengio, Y., Lamblin, P., Popovici, D. & Larochelle, H. (2006). Greedy layer-wise training of deep networks. *Advances in Neural Information Processing Systems, 19*

Bengio, Y. & LeCun, Y. (2007). Scaling learning algorithms towards AI. *Large-scale kernel machines, 34* (5), 1-41.

Berger, P.L. & Luckmann, T. (1966). *The social construction of reality: A teatrise in the sociology of knowledge*, 1st ed. Doubleday, Garden city.

Bernstein, F., Modaresi, S. & Sauré, D. (2019). A dynamic clustering approach to data-driven assortment personalization. *Management Science, 65* (5), 2095-2115.

Bicheno, J. & Holweg, M. (2000). *The lean toolbox: A handbook for lean transformation*. PICSIE books Buckingham,

Blackler, F. (1995). Knowledge, knowledge work and organizations: An overview and interpretation. *Organization Studies, 16* (6), 1021-1046.

Bloom, N., Sadun, R. & Van Reenen, J. (2012). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review, 102* (1), 167-201.

Bloom, N., Sadun, R. & Van Reenen, J. (2016). *Management as a technology?* National Bureau of Economic Research, Cambridge, Mass.

Braverman, H. (1974). *Labor and monopoly capital : the degradation of work in the twentieth century*. Monthly Review Press, London.

Bresnahan, T.F. (1999). Computerisation and wage dispersion: an analytical reinterpretation. *The economic journal, 109* (456), 390-415.

Brown, J.S. & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science, 2* (1), 40-57.

Browning, T.R. & Heath, R.D. (2009). Reconceptualizing the effects of lean on production costs with evidence from the F-22 program. *Journal of Operations Management, 27* (1), 23-44.

Brynjolfsson, E. & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science, 358* (6370), 1530-1534.

Brynjolfsson, E., Rock, D. & Syverson, C. (2019). 'Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics', in: Agrawal, A. (Ed), *The Economics of Artificial Intelligence*. University of Chicago Press, Chicago, pp. 23-60.

Burgess, T.F. (1994). Making the leap to agility: Defining and achieving agile manufacturing through business process redesign and business network redesign. *International Journal of Operations & Production Management, 14* (11), 23-34.

Campbell, M., Hoane Jr, A.J. & Hsu, F.-h. (2002). Deep blue. *Artificial intelligence, 134* (1-2), 57-83.

Chandler Jr, A.D. (1977). *The visible hand*. Belknap Press of Harvard University Press, Cambridge, Mass.

Charles, K.K., Hurst, E. & Notowidigdo, M. (2013). Manufacturing decline, housing booms, and non-employment. *NBER Working Paper Series*, 18949.

Choo, C.W. (2005). *The Knowing Organization*. New York: Oxford University Press, New York.

Choudhury, P., Starr, E. & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal, 41* (8), 1381-1411.

Collins, J.C. (2001). *Good to great: why some companies make the leap… and others don't*. Harper Business, New York.

Collins, S.T. & Browning, T.R. (2019). It worked there, so it should work here: Sustaining change while improving product development processes. *Journal of Operations Management, 65* (3), 216-241.

Curchod, C., Patriotta, G., Cohen, L. & Neysen, N. (2020). Working for an algorithm: Power asymmetries and agency in online work settings. *Administrative Science Quarterly, 65* (3), 644-676.

Cusumano, M.A. (1985). *The Japanese automobile industry: Technology and management at Nissan and Toyota*. BRILL, Leiden.

Cusumano, M.A., Holweg, M., Howell, J., Netland, T., Shah, R., Shook, J., Ward, P. & Womack, J. (2021). Commentaries on "The Lenses of Lean". *Journal of Operations Management, 67* (5), 627-639.

Davenport, T.H. (1993). *Process innovation : reengineering work through information technology*. Harvard Business School Press, Boston, Mass.

Davenport, T.H. (2005). *Thinking for a living: how to get better performances and results from knowledge workers*. Harvard Business Press, Boston, Mass.

Davenport, T.H. & Brain, D. (2018). Before automating your company's processes, find ways to improve them. *Harvard Business Review, 13 June*

Davenport, T.H. & Prusak, L. (1998). *Working knowledge : how organizations manage what they know*. Harvard Business School Press, Boston, Mass.

Davenport, T.H. & Short, J.E. (1990). The new industrial engineering: information technology and business process redesign. *Sloan Management Review, 31* (4), 11.

Dierickx, I. & Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science, 35* (12), 1504-1511.

Dietvorst, B.J., Simmons, J.P. & Massey, C. (2015). Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144* (1), 114.

Ditzler, G., Roveri, M., Alippi, C. & Polikar, R. (2015). Learning in nonstationary environments: A survey. *IEEE Computational Intelligence Magazine, 10* (4), 12-25.

Drucker, P.F. (1991). *The new productivity challenge*, Harvard Business Review, pp. 69-79.

Drucker, P.F. (2007). *Management challenges for the 21st century*. Routledge,

Eisenhardt, K.M. (1985). Control: Organizational and economic approaches. *Management Science, 31* (2), 134-149.

Epstein, J. & Klinkenberg, W.D. (2001). From Eliza to Internet: a brief history of computerized assessment. *Computers in human behavior, 17* (3), 295-314.

Evans, J.R. & Lindsay, W.M. (1989). *The management and control of quality*. West Publishing Co, St. Paul, Minn.

Faraj, S., Pachidi, S. & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization, 28* (1), 62-70.

Fayol, H. (1916). *Administration industrielle et générale: prévoyance, organisation, commandement, coordination, contrôle*. Dunod, Paris.

Feigenbaum, E.A. (1981). *Expert systems in the 1980s*. California:

Felin, T., Foss, N.J. & Ployhart, R.E. (2015). The microfoundations movement in strategy and organization theory. *Academy of Management Annals, 9* (1), 575-632.

Felin, T. & Hesterly, W.S. (2007). The knowledge-based view, nested heterogeneity, and new value creation: Philosophical considerations on the locus of knowledge. *Academy of Management Review, 32* (1), 195-218.

Ferdows, K. & De Meyer, A. (1990). Lasting improvements in manufacturing performance: in search of a new theory. *Journal of Operations Management, 9* (2), 168-184.

Ferreira, K.J., Lee, B.H.A. & Simchi-Levi, D. (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management, 18* (1), 69-88.

Form, W. (1987). On the degradation of skills. *Annual Review of Sociology, 13* (1), 29-47.

Frey, C.B. & Osborne, M.A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114* (2017), 254-280.

Fujimoto, T. (1999). *The evolution of a manufacturing system at Toyota*. Oxford university press, Oxford.

Galbraith, J.K. (1967). *The new industrial state*. New American Library, London.

Garfinkel, H. (1984). *Studies in ethnomethodology*, Special student ed. Polity, Cambridge.

Gavetti, G. (2005). Cognition and Hierarchy: Rethinking the Microfoundations of Capabilities' Development. *Organization science (Providence, R.I.), 16* (6), 599-617.

Gilbreth, F.B. (1911). *Motion study, a method for increasing the efficiency of the workman*. New York (State): D. Van Nostrand company, 1911, New York (State).

Gilbreth, F.B. (1912). *Primer of scientific management*. D. Van Nostrand, New York (State).

Gillespie, T. (2014). The relevance of algorithms. *Media technologies: Essays on communication, materiality, and society, 167* (2014), 167.

Goldin, C. & Katz, L.F. (1998). The origins of technology-skill complementarity. *The Quarterly Journal of Economics, 113* (3), 693-732.

Goldratt, E.M. (1997). *Critical chain*. North River Press, Great Barrington, Mass.

Goldratt, E.M. & Cox, J. (1984). *The goal: excellence in manufacturing*. North River Press, Croton-on-Hudson, N.Y.

Gomez-Uribe, C.A. & Hunt, N. (2015). The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS), 6* (4), 1-19.

Goodfellow, I., Bengio, Y. & Courville, A. (2016). *Deep learning*. The MIT Press, Cambridge, Massachusetts.

Gopal, A. & Prasad, P. (2000). Understanding GDSS in symbolic context: Shifting the focus from technology to interaction. *MIS Quarterly, 24* (3), 509-546.

Graetz, G. & Michaels, G. (2018). Robots at work. *Review of Economics and Statistics, 100* (5), 753-768.

Grant, R.M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal, 17* (Winter Special Issue), 109-122.

Gupta, S., Sharma, M. & Sunder M, V. (2016). Lean services: a systematic review. *International Journal of productivity and performance management, 65* (8), 1025-1056.

Hadjimichael, D. & Tsoukas, H. (2019). Toward a better understanding of tacit knowledge in organizations: Taking stock and moving forward. *Academy of Management Annals, 13* (2), 672-703.

Hafermalz, E. & Huysman, M. (2022). Please Explain: Key Questions for Explainable AI research from an Organizational perspective. *Morals & Machines, 1* (2), 10-23.

Hall, R.W. (1983). *Zero inventories*. Business One Irwin, Homewood, Ill.

Hammer, M. & Champy, J. (1993). *Reengineering the corporation: Manifesto for business revolution*. HarperBusiness, New York.

Harrison, B. (1994). *Lean and mean: The changing landscape of corporate power in the age of flexibility*. BasicBooks, New York.

Hayes, R.H. & Wheelwright, S.C. (1979). The dynamics of process-product life cycles. *Harvard Business Review, 57* (2), 127-136.

Heaton, J.B., Polson, N. & Witte, J.H. (2017). Rejoinder to 'deep learning for finance: Deep portfolios'. *Applied Stochastic Models in Business and Industry, 33* (1), 19-21.

Hebb, D.O. (2005). *The organization of behavior: A neuropsychological theory*. Psychology press,

Hedlund, G. (1994). A model of knowledge management and the N-form corporation. *Strategic Management Journal, 15* (S2), 73-90.

Hines, P., Holwe, M. & Rich, N. (2004). Learning to evolve: A review of contemporary lean thinking. *International Journal of Operations & Production Management, 24* (10), 994-1011.

Hinton, G.E., Osindero, S. & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural computation, 18* (7), 1527-1554.

Hofer, V. & Krempl, G. (2013). Drift mining in data: A framework for addressing drift in classification. *Computational Statistics & Data Analysis, 57* (1), 377-391.

Hollebeek, L.D., Conduit, J., Sweeney, J., Soutar, G., Karpen, I.O., Jarvis, W. & Chen, T. (2016). Epilogue to the special issue and reflections on the future of engagement research. *Journal of Marketing Management, 32* (5-6), 586-594.

Holweg, M. (2007). The genealogy of lean production. *Journal of Operations Management, 25* (2), 420-437.

Holweg, M., Davies, J., De Meyer, A., Lawson, B. & Schmenner, R.W. (2018a). *Process theory: The principles of operations management*. Oxford University Press, Oxford.

Holweg, M. & Maylor, H. (2018). Lean leadership in major projects: from "predict and provide" to "predict and prevent". *International Journal of Operations & Production Management, 38* (6), 1368-1386.

Holweg, M., Staats, B.R. & Upton, D.M. (2018b). Making process improvements stick. *Kenan Institute of Private Enterprise Research Paper No 18-22,*

Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L.H. & Aerts, H.J. (2018). Artificial intelligence in radiology. *Nature Reviews Cancer, 18* (8), 500-510.

Hutchby, I. (2001). Technologies, texts and affordances. *Sociology, 35* (2), 441-456.

Ingram, P. & Baum, J.A. (1997). Opportunity and constraint: Organization's learning from the operating and competitive experience industries. *Strategic Management Journal, 18* (S1), 75-98.

Inman, R.A., Sale, R.S., Green Jr, K.W. & Whitten, D. (2011). Agile manufacturing: relation to JIT, operational performance and firm performance. *Journal of Operations Management, 29* (4), 343-355.

Jaimovich, N. & Siu, H.E. (2012). *Job polarization and jobless recoveries.*

Jones, M. (2014). A matter of life and death. *MIS Quarterly, 38* (3), 895-A896.

Kalleberg, A.L. (2009). Precarious work, insecure workers: Employment relations in transition. *American Sociological Review, 74* (1), 1-22.

Kaplan, S., Milde, J. & Cowan, R.S. (2017). Symbiont practices in boundary spanning: Bridging the cognitive and political divides in interdisciplinary research. *Academy of Management Journal, 60* (4), 1387-1414.

Kellenbenz, H. (1974). Technology in the age of the scientific revolution, 1500-1700. *The Fontana Economic History of Europe, 2*, 177-272.

Kellogg, K.C., Valentine, M.A. & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals, 14* (1), 366-410.

Kieser, A. (1989). Organizational, institutional, and societal evolution: Medieval craft guilds and the genesis of formal organizations. *Administrative Science Quarterly, 34* (4), 540-564.

Knights, D., Murray, F. & Willmott, H. (1993). Networking as knowledge work: A study of strategic interorganizational development in the financial services industry. *Journal of Management Studies, 30* (6), 975-995.

Kogut, B. & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science, 3* (3), 383-397.

Krakowski, S., Luger, J. & Raisch, S. (2022). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal, 44* (6), 1425-1452.

Krishnamurthy, R. & Yauch, C.A. (2007). Leagile manufacturing: a proposed corporate infrastructure. *International Journal of Operations & Production Management, 27* (6), 588-604.

Kroezen, J., Ravasi, D., Sasaki, I., Żebrowska, M. & Suddaby, R. (2021). Configurations of craft: Alternative models for organizing work. *Academy of Management Annals, 15* (2), 502-536.

Kwak, Y.H. & Anbari, F.T. (2006). Benefits, obstacles, and future of six sigma approach. *Technovation, 26* (5-6), 708-715.

Lacity, M., Willcocks, L.P. & Craig, A. (2016). Robotic process automation at Telefonica O2. *MIS Quarterly Executive, 15* (1), 21-35.

Landes, D.S. (1969). *The unbound Prometheus : technological change and industrial development in Western Europe from 1750 to the present*. Cambridge University Press, Cambridge.

Lawler III, E.E. (1996). *From the ground up: Six principles for building the new logic corporation*. Jossey-Bass Publishers, San Francisco.

Lebovitz, S., Levina, N. & Lifshitz-Assaf, H. (2021). Is AI ground truth really "true"? The dangers of training and evaluating AI tools based on experts' know-what. *Management Information Systems Quarterly, 45* (3), 1501-1526.

Lebovitz, S., Lifshitz-Assaf, H. & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science, 33* (1), 126-148.

Lei, D., Hitt, M.A. & Bettis, R. (1996). Dynamic core competences through meta-learning and strategic context. *Journal of Management, 22* (4), 549-569.

Lenat, D.B., Guha, R.V., Pittman, K., Pratt, D. & Shepherd, M. (1990). Cyc: toward programs with common sense. *Communications of the ACM, 33* (8), 30-49.

Leonard-Barton, D. (1998). *Wellsprings of knowledge : building and sustaining the sources of innovation*. Harvard Business School, Boston, Mass.

Leonardi, P.M. & Barley, S.R. (2010). What's under construction here? Social action, materiality, and power in constructivist studies of technology and organizing. *Academy of Management Annals, 4* (1), 1-51.

Leonardi, P.M. & Vaast, E. (2017). Social media and their affordances for organizing: A review and agenda for research. *Academy of Management Annals, 11* (1), 150-188.

Levitt, B. & March, J.G. (1988). Organizational learning. *Annual Review of Sociology, 14* (1), 319-338.

Lighthill, J., Sutherland, N.S., Needham, R.M., Longuet-Higgins, H.C. & Michie, D. (1973). *Artificial intelligence: a general survey.*

Lou, B. & Wu, L. (2021). AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms. *MIS Quarterly, 45* (3), 1451-1482.

Macduffie, J.P. (1995). Human resource bundles and manufacturing performance: Organizational logic and flexible production systems in the world auto industry. *Industrial & Labor Relations Review, 48* (2), 197.

MacIntyre, A.C. (1981). *After virtue : a study in moral theory*, American ed. University of Notre Dame Press, Note Dame, Ind.

Maitlis, S. & Christianson, M. (2014). Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals, 8* (1), 57-125.

Marx, K. (1933). *Capital*. J.M. Dent, London.

McAfee, A. & Brynjolfsson, E. (2012). Big data: The management revolution. *Harvard Business Review, 90* (10), 60-68.

McCorduck, P. (1979). *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. W.H. Freeman, San Francisco.

McCulloch, W.S. & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics, 5*, 115-133.

McGahan, A.M. & Porter, M.E. (1997). How much does industry matter, really? *Strategic Management Journal, 18* (S1), 15-30.

McLachlin, R. (1997). Management initiatives and just-in-time manufacturing. *Journal of Operations Management, 15* (4), 271-292.

Miers, D., Kerremans, M., Ray, S. & Tornbohm, C. (2019). *Gartner Magic Quadrant for Robotic Process Automation Software*. Research, G.

Mintzberg, H. (1994). Rethinking strategic planning part I: Pitfalls and fallacies. *Long Range Planning, 27* (3), 12-21.

Mitchell, T.M. (1997). *Machine learning*, International ed. McGraw-Hill, London.

Mokyr, J. (1992). *The lever of riches: Technological creativity and economic progress*. Oxford University Press,

Mokyr, J. (1998). The political economy of technological change. *Technological revolutions in Europe*, 39-64.

Morales, D.R. & Wang, J. (2010). Forecasting cancellation rates for services booking revenue management using data mining. *European Journal of Operational Research, 202* (2), 554-562.

Nagel, R. & Dove, R. (1991). 21st Century Manufacturing Enterprise Strategy: An Industry Led View of Agile Manufacturing. *Vol. I and II, Iacocca Institute, Lehigh University, November,*

Narasimhan, R., Swink, M. & Kim, S.W. (2006). Disentangling leanness and agility: an empirical investigation. *Journal of Operations Management, 24* (5), 440-457.

Nelson, A., Anthony, C. & Tripsas, M. (2023). "If I Could Turn Back Time": Occupational Dynamics, Technology Trajectories, and the Reemergence of the Analog Music Synthesizer. *Administrative Science Quarterly*, 00018392231163178.

Nelson, D. (1975). *Managers and workers : origins of the new factory system in the United States, 1880-1920*. University of Wisconsin Press, Madison.

Newell, A. & Simon, H. (1956). The logic theory machine--A complex information processing system. *IRE Transactions on information theory, 2* (3), 61-79.

Nonaka, I. & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press, Oxford.

Ōno, T. (1988). *Toyota production system: Beyond large-scale production*. Productivity Press, Cambridge, Mass.

Orlikowski, W.J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization Science, 3* (3), 398-427.

Orlikowski, W.J. & Gash, D.C. (1994). Technological frames: making sense of information technology in organizations. *ACM Transactions on Information Systems (TOIS), 12* (2), 174-207.

Orlikowski, W.J. & Scott, S.V. (2008). 10 sociomateriality: challenging the separation of technology, work and organization. *Academy of Management Annals, 2* (1), 433-474.

Orlikowski, W.J. & Scott, S.V. (2015a). The algorithm and the crowd. *MIS Quarterly, 39* (1), 201-216.

Orlikowski, W.J. & Scott, S.V. (2015b). Exploring material-discursive practices. *Journal of Management Studies, 52* (5), 697-705.

Orr, J.E. (1996). *Talking about machines: An ethnography of a modern job*. ILR Press, London.

Osterman, P. (1994). How Common is Workplace Transformation and Who Adopts it? *Industrial & Labor Relations Review, 47* (2), 173.

Ouchi, W.G. (1979). A conceptual framework for the design of organizational control mechanisms. *Management Science, 25* (9), 833-848.

Pachidi, S., Berends, H., Faraj, S. & Huysman, M. (2021). Make way for the algorithms: Symbolic actions and change in a regime of knowing. *Organization Science, 32* (1), 18-41.

Pande, P., Neuman, R. & Cavanagh, R. (2001). *The Six Sigma way team fieldbook: An implementation guide for process improvement teams*. McGraw Hill Professional,

Pinch, T.J. & Bijker, W.E. (1984). The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social studies of science, 14* (3), 399-441.

Pine, K.H. & Mazmanian, M. (2017). Artful and contorted coordinating: The ramifications of imposing formal logics of task jurisdiction on situated practice. *Academy of Management Journal, 60* (2), 720-742.

Polanyi, M. (1962). *Personal knowledge : towards a post-critical philosophy*. Routledge & Kegan Paul, London.

Polanyi, M. (1967). *The tacit dimension*. Routledge & Kegan Paul, London.

Poppendieck, M. & Poppendieck, T. (2003). *Lean software development: an agile toolkit*. Addison-Wesley,

Pratt, M.G., Rockmann, K.W. & Kaufmann, J.B. (2006). Constructing professional identity: The role of work and identity learning cycles in the customization of identity among medical residents. *Academy of Management Journal, 49* (2), 235-262.

Radnor, Z.J., Holweg, M. & Waring, J. (2012). Lean in healthcare: the unfilled promise? *Social science & medicine, 74* (3), 364-371.

Raisch, S. & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review, 46* (1), 192-210.

Ramesh, G. & Devadasan, S. (2007). Literature review on the agile manufacturing criteria. *Journal of Manufacturing Technology Management, 18* (2), 182-201.

Reyes, A. & Zarama, R. (1998). The process of embodying distinctions—A reconstruction of the process of learning. *Cybernetics & Human Knowing, 5* (3), 19-33.

Ries, E. (2011). *The lean startup: How constant innovation creates radically successful businesses*. Portfolio Penguin, London, New York.

Rosso, B.D., Dekas, K.H. & Wrzesniewski, A. (2010). On the meaning of work: A theoretical integration and review. *Research in organizational behavior, 30*, 91-127.

Rucker, R. (2000). Citibank increases customer loyalty with defect-free processes. *the Journal for Quality and Participation, 23* (4), 32.

Rumelhart, D.E., Hinton, G.E. & Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature, 323* (6088), 533-536.

Rumelt, R.P. (1991). How much does industry matter? *Strategic Management Journal, 12* (3), 167-185.

Sako, M., Qian, M. & Attolini, J. (2022). Future of professional work: evidence from legal jobs in Britain and the United States. *Journal of Professions and Organization, 9* (2), 143-169.

Sampson, S.E. (2021). A strategic framework for task automation in professional services. *Journal of Service Research, 24* (1), 122-140.

Schmenner, R.W. (2012). *Getting and staying productive: applying swift, even flow to practice*. Cambridge University Press,

Schmenner, R.W. & Swink, M.L. (1998). On theory in operations management. *Journal of Operations Management, 17* (1), 97-113.

Schonberger, R.J. (1982). Some observations on the advantages and implementation issues of just-in-time production systems. *Journal of Operations Management, 3* (1), 1-11.

Schonberger, R.J. (1987). *World class manufacturing casebook : implementing JIT and TQC*. Free Press ; Collier Macmillan, New York : London.

Schroeder, R.G., Linderman, K., Liedtke, C. & Choo, A.S. (2008). Six Sigma: Definition and underlying theory. *Journal of Operations Management, 26* (4), 536-554.

Scott, S.V. & Orlikowski, W.J. (2012). Reconfiguring relations of accountability: Materialization of social media in the travel sector. *Accounting, organizations and society, 37* (1), 26-40.

Sennett, R. (2008). *The craftsman*. Yale University Press,

Senoner, J., Netland, T. & Feuerriegel, S. (2022). Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science, 68* (8), 5704-5723.

Sergeeva, A.V., Faraj, S. & Huysman, M. (2020). Losing touch: an embodiment perspective on coordination in robotic surgery. *Organization Science, 31* (5), 1248-1271.

Shah, R. & Ward, P.T. (2003). Lean manufacturing: context, practice bundles, and performance. *Journal of Operations Management, 21* (2), 129-149.

Shah, R. & Ward, P.T. (2007). Defining and developing measures of lean production. *Journal of Operations Management, 25* (4), 785-805.

Shewhart, W.A. (1926). Quality control charts. *The Bell System technical journal, 5* (4), 593-603.

Shewhart, W.A. & Deming, W.E. (1939). *Statistical method from the viewpoint of quality control*. The Graduate School, The Department of Agriculture, Washington.

Simon, H.A. (1995). Artificial intelligence: an empirical science. *Artificial intelligence, 77* (1), 95-127.

Sodhi, M.S., Seyedghorban, Z., Tahernejad, H. & Samson, D. (2022). Why emerging supply chain technologies initially disappoint: Blockchain, IoT, and AI. *Production and Operations Management, 31* (6), 2517-2537.

Sorge, A. (1991). Strategic fit and the societal effect: Interpreting cross-national comparisons of technology, organization and human resources. *Organization Studies, 12* (2), 161-190.

Spreitzer, G.M., Cameron, L. & Garrett, L. (2017). Alternative work arrangements: Two images of the new world of work. *Annual Review of Organizational Psychology and Organizational Behavior, 4*, 473-499.

Staats, B.R., Brunner, D.J. & Upton, D.M. (2011). Lean principles, learning, and knowledge work: Evidence from a software services provider. *Journal of Operations Management, 29* (5), 376-390.

Starbuck, W.H. (1992). Learning by knowledge-intensive firms. *Journal of Management Studies, 29* (6), 713-740.

Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M. & Teller, A. (2022). Artificial Intelligence and Life in 2030: The One Hundred Year Study on Artificial Intelligence.

Suchman, L.A. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge university press,

Susskind, D. (2020). *A world without work: Technology, automation and how we should respond*. Penguin UK,

Susskind, R.E. & Susskind, D. (2016). *The future of the professions: How technology will transform the work of human experts*. Oxford University Press, Oxford.

Taner, M.T., Sezen, B. & Antony, J. (2007). An overview of six sigma applications in healthcare industry. *International Journal of Health Care Quality Assurance,*

Taylor, F.W. (1911). *The principles of scientific management*. Harper & Brothers,

Teece, D.J. (2007). Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal, 28* (13), 1319-1350.

Thomerson, L.D. (2001). *Journey for excellence: Kentucky's commonwealth health corporation adopts six sigma approach*, ASQ World Conference on Quality and Improvement Proceedings. American Society for Quality, p. 152.

Tsoukas, H. (2005). *Complex knowledge: Studies in organizational epistemology*. Oxford University Press, Oxford.

Tsoukas, H. (2012). 'How should we understand tacit knowledge? A phenomenological view', *Handbook of organizational learning and knowledge management*. John Wiley & Sons, Inc, Hoboken, NJ, USA, pp. 453-476.

van den Broek, E., Sergeeva, A. & Huysman, M. (2021). When the machine meets the expert: An ethnography of developing ai for hiring. *MIS Quarterly, 45* (3)

Van Maanen, J. & Barley, S.R. (1982). *Occupational communities: Culture and control in organizations*.

Vanderschueren, T., Boute, R., Verdonck, T., Baesens, B. & Verbeke, W. (2023). Optimizing the preventive maintenance frequency with causal machine learning. *International Journal of Production Economics, 258*, 108798.

Vartak, M. (2022). How to scale AI in your organization. *Harvard Business Review,*

Vickers, G. (1965). The art of judgment: A study of policy making.

Weber, M. (1930). *The Protestant ethic and the spirit of capitalism*. Unwin University Books, London.

Weick, K.E. (1988). Enacted sensemaking in crisis situations. *Journal of Management Studies, 25* (4), 305-317.

Weick, K.E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly, 38* (4), 628-652.

Wenger, E. (1999). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press, Cambridge, UK.

White, R.E., Pearson, J.N. & Wilson, J.R. (1999). JIT manufacturing: a survey of implementations in small and large US manufacturers. *Management Science, 45* (1), 1-15.

Wiig, K.M. (1997). Knowledge management: an introduction and perspective. *Journal of Knowledge Management, 1* (1), 6-14.

Willcocks, L. (2020). Robo-Apocalypse cancelled? Reframing the automation and future of work debate. *Journal of Information Technology, 35* (4), 286-302.

Winter, S.G. (2003). Understanding dynamic capabilities. *Strategic Management Journal, 24* (10), 991-995.

Winter, S.G. (2009). 'Knowledge and competence as strategic assets', *The strategic management of intellectual capital*. Routledge, pp. 165-187.

Winter, S.G. & Nelson, R.R. (1982). An evolutionary theory of economic change. *University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship,*

Wittgenstein, L. (1953). *Philosophical investigations*. Macmillan, New York.

Womack, J.P. & Jones, D.T. (1996). *Lean thinking: Banish waste and create wealth in your corporation*. Simon & Schuster, New York, NY.

Womack, J.P., Jones, D.T. & Roos, D. (1990). *The machine that changed the world*. Rawson Associates ; Collier Macmillan Canada ; Maxwell Macmillan International, Oxford.

Wooldridge, M. (2020). Artificial Intelligence requires more than deep learning-but what, exactly? *Artificial intelligence, 289*, 103386.

Yusuf, Y.Y., Sarhadi, M. & Gunasekaran, A. (1999). Agile manufacturing:: The drivers, concepts and attributes. *International Journal of Production Economics, 62* (1-2), 33-43.

Zetka, J.R. (2003). *Surgeons and the Scope*. Cornell University Press,

Chapter 3

Learning to (be told to) count: Implementing algorithmic recommendations in retail operations

Abstract

In this paper, we study the interplay between humans and algorithms. Specifically, we examine whether algorithmic technologies can realise improvements when the efficacy of an algorithm is dependent on the level of adherence to its recommendations. We study this problem in the context of inventory record inaccuracy, seeking to minimise mismatches between physical and recorded stock levels. We first developed a deep learning algorithm that recommends product counts to retailers based on the product's likelihood of being in error. Conducting a field experiment on 450 large stores of a leading grocery retailer, we found that our deep learning model provides a 26.07% relative improvement in identifying and correcting inventory record errors that affect future reordering decisions, which in turn boosts sales and reduces wastage. At the same time, we also found that store workers enacted fewer product counts recommended by the deep learning algorithm compared to the previous heuristics-based system. In our subsequent analysis, we identify three balancing mechanisms that moderate the performance of algorithms that require human enactment of algorithmic recommendations: representation, degradation and aversion. More generally, our study highlights how generic algorithms, like deep learning, can be applied successfully to improve performance over problems that require humans and algorithms to work together to perform a task, yet also, identify the wider challenges that emerge when deploying those algorithms.

***Keywords*:** Knowledge work, process improvement, automation, machine learning, robotic process automation.

3.1 Introduction

Management researchers and practitioners alike have begun to explore the wide-ranging opportunities for algorithmic technologies to improve operational decision-making (Bastani et al., 2020; Brynjolfsson and McAfee, 2017; Chou et al., 2022). While algorithmic methods generally show compromise to improve predictions, only recently attention has been devoted to the question of how the realisation of these improvements is affected by the human actors who interface with the algorithms. Algorithmic technologies are generic computer programs that can emulate human decision-making and work (Kellogg et al., 2020); they draw insights from data and can be used in a wide range of operational decisions, such as those involved in forecasting demand (Ferreira et al., 2016), managing quality (Senoner et al., 2022), and personalising assortment (Bernstein et al., 2019). Recently, the universal benefit of algorithmic insights has been contested: some authors, for example, have expressed concerns about how algorithms may subsume key learning practices (Anthony, 2021) or obscure ways of organisational control (Kellogg et al., 2020). These studies highlight the fact that algorithms are only one ingredient in organisational work, and that operational performance is an ongoing accomplishment dependent on the interplay between organisational actors and algorithms. The notion of the interplay, thus, poses a novel challenge for operations management: how can organisations realise the promised improvements in contexts where the efficacy of an algorithm depends on the level of adherence to its recommendations?

In this paper, we study the implementation of an algorithmic technology – a deep learning algorithm – to reduce inventory record inaccuracy (IRI) in a retail setting. Deep learning algorithms are a type of machine learning algorithms that 'learn' to solve a problem using a multilayer artificial neural network architecture (Chollet, 2018). In our context, the deep learning algorithm produced insights to deal with IRI by identifying

mismatches between physical and recorded stock that distort the information retailers use to make production planning and control decisions (Chuang et al., 2022; DeHoratius and Raman, 2008; Iglehart and Morey, 1972). Despite most retailers having invested in automatic replenishment systems (Van Donselaar et al., 2010), the work to control IRI remains distributed and relies on a collaboration of operations analysts and local supermarket employees via technology to identify and correct inaccuracies in inventory. As such, IRI is a classical example of a problem in which the interplay between organisational actors and algorithms affects operational performance. Moreover, given that IRI is a long-established well-researched problem in operations management, the case represents a good test for the applicability of algorithms to create insights for operational decision-making. We use this context to situate our main research question, whether (and under what circumstances) algorithmic technologies can improve operational performance in problems that require human-algorithm collaboration.

To answer this question, we conducted a randomized field experiment involving 450 stores at a large grocery retailer, hereafter called 'Superstore'. We first develop a deep learning-enabled decision model by combining a recommender system with the cycle counting routine commonly used for auditing and correcting inventory records. In essence, our deep learning model uses historical data (240 variables including e.g., point-of-sale, inventory, and store-level data) to recommend products to be counted by supermarket workers. To measure the impact of our model on IRI, we ran a six-week randomized controlled trial at Superstore. Stores in the treatment group ($n_t = 150$) received recommendations on products to count from the deep learning algorithm, and stores in the control group ($n_c = 300$) received recommendations from operations analysts in the same way as they had been prior to the experiment. Staff at the stores were not aware of the group they were in, and we maintained the same format for

recommendations in both groups. We then used this experiment to estimate the across-store variation for multiple key performance indicators including inventory integrity, sales, food waste and product availability.

This study contributes to operations management theory by expanding our understanding of how algorithmic technologies affect operational performance in contexts that require human-algorithm collaboration. Our findings show that stores using the deep learning decision model were on average 26.07% better at identifying IRI than stores using traditional operations analysts' heuristics. This, in turn, improved operational performance by boosting sales and decreasing waste. However, we also noted a decrease in adherence to algorithmic recommendations and that the performance improvement achieved was only 46.37% of that expected. We explain these results by theorizing three balancing mechanisms: representation, degradation, and aversion. We argue that deep learning improves a process by representing problems more accurately. But these new methods bring about new challenges that need to be considered when applying them. Errors and biases may creep into the model, degrading the quality of the predictions; and employees may avert algorithmic recommendations as a natural adaptation process preventing the company from realising expected benefits. These findings suggest that to best augment operations with algorithmic technologies, it is important to be aware that algorithms are not self-standing entities, but that they operate with humans in an ensemble with different dynamics at play.

3.2 Theoretical foundations

3.2.1 Using algorithms to augment operational performance

Access to exponentially growing data combined with progress in novel algorithmic technologies has greatly improved organisations' ability to make predictions across a wide range of operations management areas (Bastani et al., 2020; Mišić and Perakis, 2020; van der Haar et al., 2023). Particularly, machine learning algorithms have received a great deal of attention because of their ability to learn complex patterns from historical data (Brynjolfsson and Mitchell, 2017; Csaszar and Steinberger, 2022; Raisch and Krakowski, 2021). Much of this work treats algorithms as a general-purpose technology (Brynjolfsson and McAfee, 2017) that can support decision-making for a wide range of organisational problems. These problems are either framed as a regression (i.e., predicting continuous variables), classification (i.e., predicting discrete categories) or ranking task (i.e., predicting policies that maximise a reward function). It is now widely acknowledged that machine learning algorithms are an effective tool for key operations problems, such as demand forecasting (Ferreira et al., 2016), revenue management (Morales and Wang, 2010), dynamic assortment personalisation (Bernstein et al., 2019), and quality improvement (Senoner et al., 2022).

Concerning enhancing operational performance, an emerging research agenda has begun to show improvements from using algorithms to automate knowledge tasks – work components that require human thinking. In general, improvements stem from increased resource efficiency. UPS, for example, saved circa $350 million annually from implementing algorithms that create optimal routes for parcel collection and delivery (Holland et al., 2017). Improvement can also result from reducing the time workers take to perform a task. For instance, the Alibaba group lowered packing time by 4.5% by following the algorithm's bin packing recommendations (Sun et al., 2021). While in

others, improvement is achieved by minimising workers' load. To illustrate, algorithms for targeted automation of order decisions free planners' time for other value-added activities (Imdahl et al., 2021).

In addition to improving resource efficiency by automating knowledge tasks, some algorithms may improve work by increasing organisational information processing, generating insights previously unavailable to decision-makers (Murray et al., 2021). Critical to this position is the view humans are subject to "bounded rationality" – a state of incomplete knowledge and anticipation of consequences that leads to making satisfactory rather than optimal decisions - and "myopic search" – the process of finding nearby solutions rather than afar for a given problem (Cyert and March, 1963; Gavetti et al., 2012; March et al., 1958; Simon, 1947). Unlike humans, emergent algorithms synthesise patterns in vast amounts of data quickly, for example, through an inductive learning process (Choudhury et al., 2021). This capability enables organisations to create new decision-making structures (Shrestha et al., 2019) and mitigate the shortcomings of human decision-making; it results in new ways of finding solutions and selecting actions that are closer to being optimal, rather than merely satisfactory (Murray et al., 2021).

Deploying algorithms may improve organisational information processing even further by enhancing the representation of a problem. A representation is a mental model, which plays an important role in focusing actors' attention on some facets of the environment (Gavetti and Levinthal, 2000) and subsequent decision-making (Simon, 1991). Human representations are often simplistic (Thagard, 2005), as they are limited by actors' inability to parse complex spatial (Porac et al., 1989), temporal and causal relationships (Weick, 1979). When organisations replace human decision-making with algorithmic predictions, more complex and holistic representations of problems may become feasible.

One set of algorithms that in particular can improve representations are multi-layered artificial neural networks, commonly known as deep learning (Chollet, 2018). These algorithms use a network architecture of connected nodes (modelled on the human neuron) to learn multiple levels of representations by composing simple (yet non-linear) modules that each transform the representation at one level into a representation at a higher level (Chollet, 2018; LeCun et al., 2015). Typical examples include convolutional neural networks for image recognition and transformers for natural language processing. Deep learning differs from other traditional algorithms such as logistic regression, random forests, and adaptive boosting classifiers, by not requiring humans to identify and transform the relevant features and relationships that affect a problem. Many algorithms may improve operational performance, but it is generally more likely that improvement emerges from those with architectures that augment problem representations, such as deep learning because they relax the limitations of humans' bounded rationality.

3.2.2 The interplay between humans and algorithms

Recently, attention has been given to the interplay between humans and algorithms. The core idea of this stream of research is that algorithms are not self-standing, independent artefacts, but rather that humans and algorithms interact in an 'ensemble' (Leonardi and Barley, 2008; Orlikowski and Scott, 2008, 2015; Scott and Orlikowski, 2014). As Glaser et al. (2021, p.20) noted: algorithms are 'constituted in wider agentic sociomaterial assemblages whose relational properties evolve and shift as they are re-enacted, with varying degrees of performativity'. Expressed in this way, the notion of the ensembles runs against the thesis of the universal benefit of algorithmic technologies and calls for understanding how the cojoined agency between humans and algorithms creates change, and in turn performance gains for organisations (Murray et al., 2021). One aspect of

particular relevance to this study is ensembles that arise as organisational actors receive recommendations from algorithms.

Researchers have noted that when organisations introduce algorithms into work, actors at the interface may react, affecting the enactment of algorithmic advice and, thus, in turn, its impact on performance (Dietvorst et al., 2015, 2018). Some reactions are affective and include actors' feelings of anxiety, vulnerability, and mistrust towards algorithms. In a study of eBay business sellers, for example, Curchod et al. (2020, p.645) noted that sellers felt they were 'part of an obscure mechanism of algorithm-generated metrics that diminishes their agency'. Emotional responses to algorithms may trigger aversion. As organisational actors perceive algorithms as a threat to their work, status, or power, they may purposefully resist algorithms to prevent change to the status quo. In some cases, humans have been found to even act symbolically and only pretend to comply (Pachidi et al., 2021). Organisational actors, in general, are less likely to enact algorithmic recommendations if they perceive the algorithm as a threat.

Aversion to algorithmic recommendations may also be triggered by organisational actors' cognitive reactions (Burton et al., 2020), as a consequence of human biases (De Baets and Harvey, 2018; Eroglu and Croxton, 2010; Oliva and Watson, 2009) and lack of understanding (Anthony, 2018, 2021; Lebovitz et al., 2022). Commonly cited examples of biases include overconfidence and anchoring bias. In the former, actors tend to be overconfident in their own knowledge and ignore algorithmic advice – even when the advice is perfectly correct (Allen and Choudhury, 2021; De Baets and Harvey, 2018; Gardner and Berry, 1995). In the latter, actors use algorithmic advice as a mental anchor and then judgementally readjust it without producing significant improvements (De Baets and Harvey, 2018). A lack of understanding of the algorithmic recommendations may also result in 'algorithm(ic) aversion' (Yeomans et al., 2019). Algorithms in general,

and deep learners in particular, are characterized by being opaque because of the difficulty of understanding the reasoning behind a given outcome (Faraj et al., 2018). Although some actors may adopt algorithm recommendations without questioning their validity (Anthony, 2021), others may experience increased uncertainty when these recommendations diverge from their judgment. As a result, actors may opt to ignore the algorithm's advice (Lebovitz et al., 2022) or replace it with their own judgment (Waardenburg et al., 2021). In either case, these cognitive reactions hamper the realisation of improvements from investments in algorithmic technologies.

In summary, these two bodies of literature point to one key question that requires more attention: whether (and under what circumstances) algorithmic technologies are able to improve operational performance in problems that require human-algorithm collaboration.

3.3 Focal problem

3.3.1 The study of inventory record inaccuracy

We consider the implementation of a deep learning algorithm to address the problem of inventory record inaccuracy, or IRI for short. In essence, IRI denotes a mismatch between the physical and recorded stock for a stock-keeping unit (SKU), which distorts the information a firm uses to make production planning and control decisions (Chuang et al., 2022; DeHoratius and Raman, 2008; Iglehart and Morey, 1972; Rinehart, 1960). IRI is a problem for manufacturers and retailers alike as both parties rely on accurate inventory records to make key operational decisions such as when and how much stock to order from a given supplier to meet consumer demand. Inaccurate information can result in firms placing orders with the wrong quantities and/or at the wrong time,

propagating information distortions upstream to suppliers and potentially triggering the bullwhip effect (Lee et al., 1997).

At the firm level, IRI threatens the firm's operational performance by causing out-of-stock events and overstocking (Chuang and Oliva, 2015; Corsten and Gruen, 2004; Kang and Gershwin, 2005; Raman et al., 2001; Rinehart, 1960). Out-of-stock events, in turn, have been shown to cause the firm to lose up to half of the intended customer's purchases (Corsten and Gruen, 2004), while a recent study of ten consumer packaged goods categories found that out-of-stock events accounted for nearly $3 billion in lost sales in the US in 2020 (Redman, 2021). A further yet less discussed consequence of IRI arises when the stock record understates product quantities, showing a balance insufficient to meet customer orders and incorrectly triggering a product purchase order. Here, reordering drives overstocking, increasing the cost of carrying inventory, labour and waste, which is in particular an issue for products affected by shelf-life and spoilage (Raman et al., 2001).

IRI is a classical operational problem with multiple causes that make it difficult to resolve in practice (e.g. Chuang and Oliva, 2015; Delaunay et al., 2007; Rinehart, 1960; Sarac et al., 2010). Some causes are procedural in nature, as errors stem from mistakes in executing operational processes, such as sales and restocking, replenishment, assortment, and correction processes (DeHoratius and Raman, 2008; Lee and Özer, 2007; Raman et al., 2001; Ton and Raman, 2010); some are epistemic, as errors result from poor master data quality used to populate inventory tracking systems (DeHoratius and Raman, 2008; Iglehart and Morey, 1972; Rinehart, 1960); while others relate to shrinkage, as errors arise from a reduction of the physical inventory due to damage, spoilage, and pilferage (Kang and Gershwin, 2005). The table in Appendix A provides a summary of the sources and root causes of IRI, as well as associated error directions.

Not surprisingly, the diverse set of root causes for IRI reduces the effectiveness of general solutions that have been proposed in the literature, such as investing in RFID technologies (Hardgrave et al., 2013; Hardgrave et al., 2011; Heese, 2007; Uçkun et al., 2008) or partnering with suppliers to exchange information and increase demand visibility (Holweg et al., 2005). Instead, in practice, most firms rely on corrective strategies such as counting routines to identify and correct stock discrepancies (Chuang, 2015; DeHoratius et al., 2008; Kang and Gershwin, 2005).

Counting routines require multiple actors and technology to work together continuously: actors either count a subset of (i.e. cycle counting) or the entire inventory (i.e. full counting) in a given location (Brooks and Wilson, 2007; Chuang et al., 2016; Iglehart and Morey, 1972; Kök and Shang, 2007). The decision of which SKUs to count generally rests with a decision-maker (usually an operations analyst) and is based on a combination of historical analyses and the analysts' tacit knowledge and expertise to identify the optimal subset of SKUs to count in a given location and day. In the case of cycle counting routines, operations analysts commonly define a set of heuristics to prescribe a list of targeted products to count (hereafter simply referred to as 'counts'). A list of counts is then sent to the actors in the store who count and update the stock records held in the system. Solutions to successfully address IRI thus require an effective collaboration between operational analysts, technology, and the store employees that enact their count requests.

3.3.2 Research context

Our research is based on the operations of a leading grocery and merchandise retailer, hereafter referred to as Superstore. Superstore operates stores in five countries across Europe, employs c.420,000 employees, and reports annual sales of £53.4bn in 2021. For our study, we focused on Superstore's operations in the United Kingdom, where the

company has 3,433 stores of various sizes, and also operates a multichannel distribution system. The current store portfolio comprises large hypermarkets ($S_H = 252$), supermarkets ($S_S = 479$), smaller stores in central locations ($S_M = 1{,}740$), and micro convenience shops ($S_C = 956$).

Superstore uses cycle counting routines to identify and correct errors in their inventory records to streamline replenishment operations and prevent out-of-stock events. To define its cycle counting routines, operations analysts at Superstore combine heuristics, root cause analysis tools, and inspection policies (see Appendix B) to predict inventory records that are most likely to be in error; then analysts compile and send a list of the products to count to each store. Given that retailers will allocate a limited budget to stores to correct their inventory, analysts filter and prioritise the count recommendations $\hat{C}$ to be sent to the stores. Once the actual count of the physical stock on the shelf and backroom has been conducted by employees in the respective stores, noting any discrepancy in the records. Once the records are updated, analysts review how many of their product recommendations were in error and propose new counts for the next inspection period. We formalise the existing decision model in the following steps:

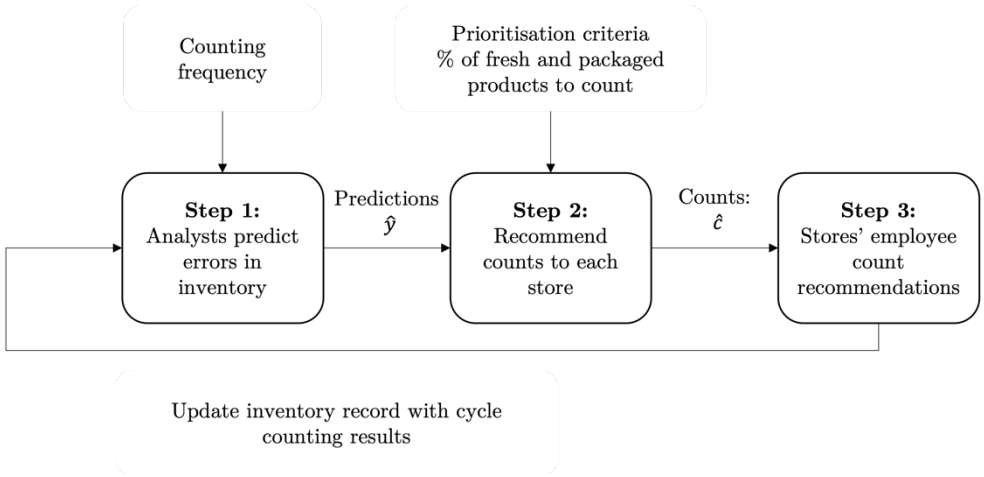


Figure 3-1 Operations analysts' decision model for correcting IRI based on periodic cycle counting routines.

The routine depicted in Figure 3-1 has been in operation at Superstore since the early 2000s, when Superstore reassigned the responsibility of managing shelves – that is ordering SKUs and replenishing shelves – from store managers to the supply chain team. This integration strategy provided the company with increased visibility on demand and inventory patterns, allowing it to make better replenishment, assortment, and promotion decisions. Reassigning store-level decisions to the supply chain team further centralised operations, which subsequently facilitated the deployment of new initiatives and technologies at Superstore.

One of these technologies is called the *Targeted Counts System* - a tool that manages the cycle counting policy and pushes product counts for each store based on a set of heuristics and analysts' recommendations. An example of these heuristics is the suspicious minimum heuristic, whereby the heuristic looks for "freezing" (Kang and Gershwin, 2005) – situations where '*the inventory management system registers a recorded inventory large enough to prevent replenishment, even though there is no physical inventory and thus no sales' (DeHoratius et al., 2008).* Thus, since the probability of continuously hitting zero sales is extremely low, the suspicious minimum heuristic looks for products with zero or no sales during a period to infer where the in-hand inventory is zero. Then it either auto-corrects the inventory records or recommends counting those products to the store. Other computer heuristics used by Superstore are summarized in Appendix B.

In 2019, Superstore also implemented an in-house random forest algorithm to predict inventory record inaccuracies, in their targeted count system. Under this new logic, operation analysts use the target count system to recommend stores with 30-70 products to count daily, based on store format. These recommendations combine the

operation's analyst suggestions based on root cause analysis, the heuristics described above, and errors identified by the random forest model. This strategy proved to be an effective way to improve inventory integrity, allowing stores to correct discrepancies in their stock quickly. At the time, this model increased sales-based availability (~0.16%), but also an increase in waste. Interested in learning more about the effects of other meta-models on improving IRI, Superstore allowed us to deploy our decision model described in section 3.4 in a randomized experiment on a portion of their state.

In addition to internal cycle counting efforts, a third-party vendor audits the full inventory every 6 to 12 months at each store; Superstore's employees run cycle counting routines and audit a subset of products daily per store.

3.3.3 Experimental setup

Superstore rolled out our deep learning-based counting system (see section 3.4) in October 2021, involving 60% of the supermarket and hypermarket stores in the United Kingdom. The goal of this experiment was to identify the effect of our meta-model on inventory accuracy and other performance indicators by comparing the post-policy change in performance on the stores deploying our meta-model against the change in performance on stores using traditional correction strategies (that is the aggregated counts suggested by the operation analysts, the automatic heuristics, and random forests). We allocated the stores into treatment ($S_{treatment} = 150$) and control groups ($S_{control} = 300$) using a random sampling strategy and concealed to the store managers information about whether they were in the treatment or control group to prevent them from adjusting their behaviour. Note that small stores were not included in the experiment, instead, the control and treatment groups were sampled from a pool of large stores as they are more complex to manage, and so are more likely to exhibit larger inventory record inaccuracies. We collected data before and after the policy was

implemented and compared the results on the following operational metrics. Our randomized sampling approach is presented in Table 3-1.

Table 3-1 Characteristics of the randomized store allocation for our experiment

Characteristics	Control group		Treatment group	
	Supermarkets	Hypermarkets	Supermarkets	Hypermarkets
No. of stores	200	100	100	50
Avg, No. of SKUs	7,358.8	11,810.7	7,554.5	12,059
(std)	(2,350.9)	(1,599.5)	(2,154.0)	(1,594.5)
Normalized stock control hours	0.60	1.00	0.62	1.00
Price Area 1	154	81	76	41
Price Area 2	29	12	15	6
Price Area 3	15	7	7	3
Price Area 4	2	0	1	0

3.4 Deep learning-based decision model

3.4.1 Model logic

We employ a four-step data-driven cycle counting decision model to improve inventory record inaccuracy, as illustrated in Figure 3-2. Similar to Senoner et al. (2022) study on the application of explainable artificial intelligence in quality management, we adopt a process flow approach to explaining our decision model. In Step 1, we use a machine learning algorithm to develop a meta-model f that leverages patterns in historical data X to estimate the probability y of a product record being an impactful error. In Step 2, the model uses unseen data to predict all inventory records in error at a given point t.

As firms allocate a limited budget to correct stock inaccuracies, in Step 3, we use prioritisation criteria to recommend each store with a set of products to count. Completing the cycle, in Step 4, store employees manually count the suggested products and update inventory records – information fed back to the model to update its beliefs and estimations. In the following, we discuss each step in detail.

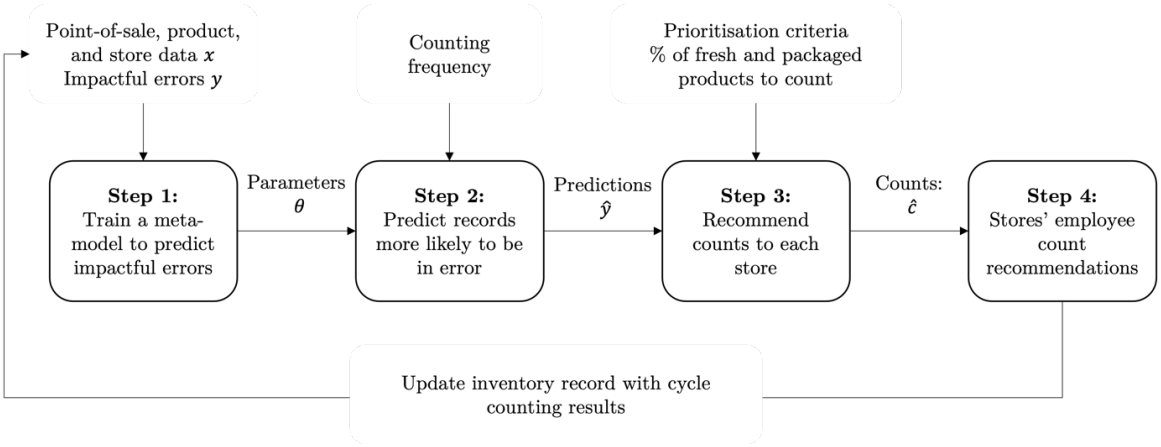


Figure 3-2 A deep learning-based decision model for the cycle counting routine in the context of grocery and merchandise retail

Step 1: Training a meta-model to predict impactful errors.

The basis of our decision model is a binary classification meta-model $f: \mathbb{R}^N \longrightarrow \mathbb{R}$ that predicts impactful errors y based on historical data x, formally, $y = f(x)$. In this step, we were concerned about learning the right meta-model f, which included considerations about the target and predictor variables, as well as underlying machine learning techniques used to map the data into predictions. In essence, the meta-model is a classification model that learns non-linear data patterns in a high-dimensional space.

First, the meta-model f is trained using historical data from inventory audits – periodical processes whereby every product at every store is counted to correct any discrepancy between the recorded and in-hand inventory. These audits contain

information about the likelihood of an error, its magnitude, and direction. We mapped these variables into the dependent variable impactful error y – a binary variable that takes 1 when an error is impactful, otherwise 0. Using data mining techniques, we augmented the training set with 51 other variables including historic point-of-sale data, inventory levels, waste, store size and product dimensions among others to make our decision models. Our theoretical feature selection was motivated by earlier work on inventory and assortment management Some features were selected by analysing key variables used in traditional cycle count policies (see Table 3-2). Others were selected based on further analysis of the drivers of inventory record inaccuracies including variables that capture the store's product variety and inventory levels, which increase the complexity of managing a store, and so the IRI (Ton and Raman, 2010); some variables such as dollar cost, it is well accepted that the higher the product's value, the lower its IRI, as managers and employees pay more attention to errors on these products. For more details on key predictor variables derived from the literature please see Table 3-2.

Table 3-2 Drivers and variables that may affect inventory record inaccuracies.

Drivers	References	Variable	Expectation
Audit Frequency	(Chuang et al., 2016; DeHoratius and Raman, 2008; Hardgrave et al., 2013)	Number of audits in a given period	The more frequent an audit, the lower the IRI
Time between audits	(Morey and Dittman, 1986)	Number of days since last audit	The more recent an audit, the lower the IRI
Dollar value/dollar volume	(DeHoratius and Raman, 2008)	Product cost	The higher the dollar value and the dollar volume of an SKU, the lower the IRI level.
Transaction frequency	(Bernard, 1985; DeHoratius and Raman, 2008; Rinehart, 1960; Ton and Raman, 2010)	Quantity sold	IRI levels increase with sales velocity

Product variety	(DeHoratius and Raman, 2008; Levy et al., 2004; Raman et al., 2001)	Store type as a proxy for product variety Inventory density as number of units found in a store selling area (DeHoratius and Raman, 2008) Product variety as 'number of merchandise categories within a store' (Levy et al., 2004)	The higher the product variety, the higher the IRI
Employee turnover	(Raman et al., 2001)	Employee turnover, managers turnover.	The higher the employee turnover, the higher the IRI
Distribution structure	(DeHoratius and Raman, 2008)	Dummy for whether the product is shipped from an owned distribution centre	
Product size	(Ton and Raman, 2010)	Product Height Product Length Product Size Product Weight	Smaller products may also be more difficult to differentiate than larger product
Storage size	(Ton and Raman, 2010)	Size of storage area	The bigger the storage area the higher the IRI

Step 2: Predicting records more likely to be in error

In this step, we are interested in predicting the inventory records more likely to be impactful errors to minimise stockouts and waste. Predictions are made usually on a daily or weekly basis, and so we apply our meta-model on unseen transactional data for the model to make predictions. While it would be convenient to recommend counting all the products that the meta-model f predicts are in error, such a solution is economically unviable. Commonly, retail stores have a limited budget to correct IRI, hence a prioritisation mechanism is necessary to decide what inventory records to audit. Considering this, in this step we provide a measure of the probability of inventory inaccuracy to base on which retailers can prioritise internal counting routines. In the case of deep neural networks, the measure of probability is given by a sigmoid function.

Step 3: Recommending counts to store's employees

Now that the model has predicted which stock records are more likely to be in error, the next step is to decide which products to prioritise for a count in each store. This prioritisation is constrained by stores' budget for cycle counting, which regulates the frequency and quantity of products that can be counted during an inventory audit. For example, a retailer may opt to count 70 products per day in hypermarkets yet only 40 products per day in supermarkets, as larger stores deal with higher product variety and thus are more likely to have higher IRI (Ton and Raman, 2010). Retailers also prioritise cycle count audits based on product categories. For example, 50% of the daily counts may be allocated to improve IRI for products in the fresh category, while the other half may be distributed among the rest of the store product categories. At the end of this step, each store receives a list of products that are more likely to present discrepancies that may impact future reordering decisions.

Step 4: Counting products and correcting stock records

The final step of our decision model involves assigning stores' employees to audit the recommended products. Employees take the list of recommended products generated by the model and visit the shelves and backroom area to count the number of units on hand for each product. After verifying if a discrepancy exists and its size, the employees update the inventory records with their findings. The data is automatically exported into the Data Warehouse, updating the information that the model would use for making predictions the following day.

3.4.2 Model's data and validation

To build our decision model, Superstore provided us with historical data on audits conducted across stores in the United Kingdom during the last year. For training the

deep learning model we decided, however, to use only data from audits conducted during the last three months, involving approximately 39,678 unique products. This decision reflected the need to capture seasonality changes that affect replenishment decisions, for example by choosing to replenish products more frequently during the Christmas period to accommodate for a desired surplus in demand. Then, we combined this information with administrative data about product characteristics, such as product category, cost, dimensions, point-of-sale, and historical inventory records; as well as store characteristics such as store format, audit frequency, and time between audits.

In total, our model accounted for 56 variables. It should be noted that while the data used in steps 1 and 2 are different, we maintain the same features. The only exception is the target variable, which is unknown at the point the predictions are made. We also collected information about the budget allocation for fresh and packaged products at the store level, which we later used for prioritising recommended counts (see Step 3). Finally, to complete the model, we collected daily information about the counts performed at the store, including the percentage of products counted after our recommendations (compliance rate), and the on-hand inventory recorded after the count (Step 4) This allowed us to track the performance of our model at different stages.

Descriptive statistics

Figure 3-3 presents the histogram of the product range by store size. On average, we find that hypermarkets in our sample manage 19,704 products, of which 3,119 are fresh, while supermarkets manage 12,915 products of which 2,419 are fresh (for more descriptive statistics please see Table 3-3). Our analysis shows that 65.38% of the inventory records in our sample are inaccurate, which is consistent with DeHoratius and Raman (2008) findings. More relevant for this study still, it shows that 13.58% of the inventory records have discrepancies that would have impacted replenishment decisions. These

discrepancies range from a minimum of 1 unit per item to a maximum of 40 units per item.

Table 3-3 SKU composition of the sample data, categorized by store format

Store Format	Average SKUs	Min SKUs	Max SKUs	% Fresh SKUs	% Packaged SKUs
Hypermarket	19,704	16,999	22,192	15.83% (1.77%)	84.17% (5.97%)
Supermarket	12,915	4,686	19,208	16.68% (5.61%)	83.32% (24.91%)

We also calculated the percentage of inventory records within various quantity discrepancy sizes (see Figure 3-3). Our analysis reveals that only a small percentage of absolute discrepancies below five units would have impacted replenishment decisions. While most non-impactful errors range between zero and five units, most impactful inventory errors vary between 6 and 20 units. The distribution also shows a long tail, as large errors are probably more easily identified and corrected.

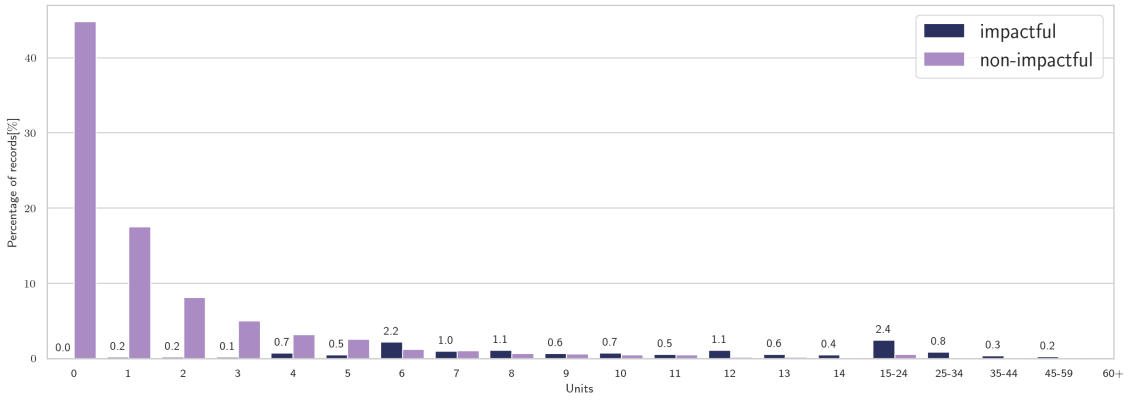


Figure 3-3 Histogram of the absolute difference between recorded and in-hand inventory

Validation strategy

To select the best meta-model to use in our experiment (Step 1), we first trained multiple deep-learning supervised classifiers with different architectures, including different combinations of neurons and hidden layers (see Table 3-4 for a list). The meta-model

learns by optimising the binary cross-entropy. Since the training examples show an imbalance between impactful and non-impactful errors (only 13.57 % of the training examples are impactful errors), we considered our problem as an instance of an imbalanced classification problem. We thus compared results augmenting the data with synthetic minority over-sampling technique (SMOTE) (Chawla et al., 2002). We then compared and selected the top-performing model based on the area under the ROC, as this metric captures the ability of the classifier to distinguish between classes, in our case, impactful and non-impactful errors. We selected a meta-model with 5 hidden layers, l00 neurons per layer, and data sampled with SMOTE.

Table 3-4 Performance comparison for deep learning meta-models

No Neurons	No Hidden Layers	Batch Size	SMOTE	ROC
200	5	100	No	0.83
100	5	100	No	0.81
100	5	100	Yes	0.91

SMOTE stands for Synthetic Minority Over-sampling Technique, which is a method used to address class imbalance in machine learning. On the other hand, ROC refers to the Receiver Operating Characteristic curve, a graphical representation of a classifier's performance, with values ranging from 0 to 1.

After the meta-model was trained, we used it to predict the probability of each stock-keeping unit being in error at each facility (see Table 3-5). Using the probability of impactful error, the model ranked the stock-keeping units and estimated whether it is likely to be in error using a threshold of 0.5.

Table 3-5 Examples of SKUs predicted to be in error

Stock-Keeping Unit	Store ID	Date	Probability of impactful error

50154381	2102	11/11/2021	0.86
85511429	2141	11/11/2021	0.86
51620087	2164	11/11/2021	0.9
88887814	2701	11/11/2021	0.86
86014613	2742	11/11/2021	0.83
62489473	2693	11/11/2021	0.84
51583337	5520	11/11/2021	0.83
88894310	6286	11/11/2021	0.84

The Probability of impactful error is a value between 0 and 1, with 1 indicating that the algorithm was fully confident that the recommended product inventory record was in error.

We now present the effect of using our model to predict discrepancies in the stores' full counting audits (static data). This is a sensible approach as the data used to train our meta-model was static; it resulted from taking a snapshot of all inventory discrepancies across a portfolio of stores during discrete events (audits). Furthermore, since full counting programmes are outsourced to third-party vendors rather than to internal employees, we would expect the counts to be more accurate, by being less susceptible to human errors or agent's strategic behaviour. Thus, it is only fair to assess the performance of our meta-model first on static data from previously unseen audits.

To that end, we monitored the performance of our model using a confusion matrix (see Table 3-6). This matrix highlights the number of discrepancies correctly classified, as well as type I and type II errors. Because our meta-model aimed at identifying discrepancies that affect replenishment decisions, and this is an imbalanced classification problem, we avoided using accuracy as a performance metric because it would not provide an appropriate evaluation of the model. Instead, we used specificity (0.69) and sensitivity (0.78).

Table 3-6 Performance of the meta-model on full counting audits

		Actual value	
		p	n
Predicted value	p'	40303 (TP)	100507 (FP)
	n'	11301 (FN)	227964 (TN)
Total		51604 (P)	140810 (N)

p and n represent the number of recommended products for which correcting their inventory records resulted in identifying impactful errors and non-impactful errors, respectively. Similarly, p' and n' represent the number of recommended products predicted to have impactful errors and non-impactful errors in their inventories, respectively.

3.4.3 Key variables

Dependent variables

A range of ways to measure stock accuracy have been proposed in the literature. Schrady (1970), for example, studied IRI in the context of the naval supply systems and documented an exhaustive set of operational metrics that describe the likelihood, severity, and cost of IRI. Some studies measure IRI using binary errors, arguing a record is in error, either if it differs from the on-hand stock (Raman et al., 2001), or if such difference exceeds a prespecified tolerance – usually 1 % (Emma, 1966; Rinehart, 1960; Sheppard and Brown, 1993). Other studies argue for measuring the severity of IRI, defining errors as the absolute difference between the recorded and physical stock (DeHoratius and Raman, 2008).

While these metrics describe the probability and magnitude of IRI, the remaining challenge is that they treat all inventory inaccuracies equally, hindering companies from prioritising records that need to be corrected sooner to ensure product availability. Addressing this shortcoming, we differentiate between 'impactful' errors – those that affect replenishment decisions – and 'non-impactful' errors - those that do not. The focus on impactful errors aims to capture the link between fixing errors, replenishment decisions and the company performance. Consider, for example, the different effects of

impactful and non-impactful errors on retailer sales. An internal analysis at superstore estimates that they lose three times more sales during out-of-stock events caused by not detecting impactful errors than by not detecting non-impactful errors. Similarly, retailer sales take longer to recover from impactful errors. We thus argue that companies should prioritise the identification and correction of impactful errors as they have a greater impact on the organisation's operational performance.

Detection of inventory record inaccuracy

We analysed the effect of deep learning on inventory accuracy using the impactful count rate (ICR) variable. This variable estimates the portion of impactful counts detected during a cycle counting routine. To illustrate, an ICR of 0.3, for example, means that the store identified 30 impactful errors in 100 SKUs counted. Unlike the store's total inventory accuracy, the ICR can be measured daily and does not require auditing every SKU in the store. Thus, the ICR is a metric well-suited to capture the stores' ability to identify and correct errors in a given period t. We calculate the impactful count rate using equation 1, with all variables defined in Table 3-7

$$\overline{ICR_t} = \sum_{s=1}^{S} \frac{IMPACTFUL_COUNTS_{st}}{S * CP_{st}} \quad (1)$$

Table 3-7 Variable definition for the impactful count metric

Variable Name	Definition
$\overline{ICR_t}$	Average impactful count rate across all stores S a given group in a day t
$IMPACTFUL_COUNTS_{st}$	Number of recommended products that were counted, and as a result of correcting their inventory records, led to a change in the subsequent purchase order for at least one case, at store s in day t
S	Number of stores in the experiment

C_{st}	Number of recommended products counted at store s in day t

Sales

Similar to our investigation of the detection of inventory record inaccuracies, we examine the effect of deep learning on sales by calculating the sales forecast error (SFE) at given time t. This metric aims to capture the cost of any stockout events that may result from poor reordering decisions caused by IRI; it measures the sales lost due to an SKU not being available in relation to the sales. One key limitation of this metric, however, is that this metric disregards the substitutability of any given SKU. We calculate the sales forecast error as the difference between forecast and actual sales for all the SKUs recommended to count at each store. We divide this by the number of total counts performed at each store at day t (C_{st}) to estimate the sales forecast error per count. The sales forecast error equation is presented below, with its variables defined in Table 3-8.

$$\overline{SFE_t} = \sum_{s=1}^{S} \frac{\sum_{c=1}^{C_s} (FORECASTED_SALES_{cst} - ACTUAL_SALES_{cst})}{S * C_{st}} \quad (2)$$

Table 3-8 Variable definitions for the sales metric

Variable Name	Definition
$\overline{SFE_t}$	Average sales forecast error across all stores S a given group in a day t
$FORECASTED_SALES_{cst}$	Predicted sales for counted product c at store s in day t
$ACTUAL_SALES_{cst}$	Actual sales for counted product c at store s in day t
S	Number of stores in the experiment
C_{st}	Number of recommended products counted at store s in day t

Food waste

Further, we measure the change in product wastage, or waste for short, as a percentage of sales. Calculating waste was central to exploring the trade-offs between sales and food waste of our decision model. Previous studies have traditionally focused on the impact of IRI on sales due to stockouts, suggesting that having excess inventory may overcome this issue. Excess inventory, however, also increases waste and inventory carrying costs. Thus, we calculate the impact of our policy also on waste, using equation 3. The variables for this equation are presented in Table 3-9.

$$\overline{WPS_t} = \sum_{s=1}^{S} \frac{WASTE_{st}}{S * SALES_{st}} \tag{3}$$

Table 3-9 Variable definitions for the waste metric

Variable Name	Definition
$\overline{WPS_t}$	Average waste per sale across all stores S in a given group in a day t
$WASTE_{st}$	Total cost of waste at store s in day t
$SALES_{st}$	Total sales at store s in day t
S	Number of stores in the experiment

Availability

Finally, we analyse the product availability using the sales-based availability ratio as a proxy. This ratio measures the probability of an item has a gap based on the frequency of sales. We calculate this metric for fast-selling and star-line products as they respond faster to demand fluctuations or are more likely to produce stockouts and impact the retailer's bottom line. Formally, sales-based availability (SBA) is estimated in the for

fast-selling products (equation 4) and for star line products (equation 5). The variables used in these equations are presented in Table 3-10.

$$\overline{\mathrm{SBA_t}} = \sum_{s=1}^{S} 1 - \frac{\sum_{f=1}^{F_s} GAPS_{fst}}{S * F_s} \tag{4}$$

Similarly, we estimate the sales-based availability of products that matter most (PMMA) as follows:

$$\overline{\mathrm{PMMA_t}} = \sum_{s=1}^{S} 1 - \frac{\sum_{l=1}^{L} GAPS_{lst}}{S * L_s} \tag{5}$$

Table 3-10 Variable definitions for sales-based availability metrics

Variable Name	Definition
$\overline{\mathrm{SBA_t}}$	Average sales-based availability for fast-selling products across all stores S in a given group in a day t
$GAPS_{fst}$	Number of gap scans for all fast-selling products at store s in a day t
$\overline{\mathrm{PMMA_t}}$	Average sales-based availability for star line products across all stores S in a given group in a day t
$GAPS_{lst}$	Number of gap scans for all-star line products at store s in a day t
F_s	Number of fast-selling products at store s
L_s	Number of star line products at store s
S	Number of stores in the experiment

3.5 Findings

3.5.1 Overall performance effect

After evaluating the impact of our model against the data on the full counting routine, we used our field experiment to evaluate the performance of stores in the treatment and control groups. We first focused on the ability to identify inventory record errors. We visually inspected the parallel trend assumption for the outcome effect estimation. Figure 3.4 compares the number of counts that resulted in detecting an impactful error for stores in the control and treatment groups before and after the treatment. The vertical line corresponds to the introduction of the machine-learning recommendation policy on $t = 0$ (October 30$^{\text{th}}$, 2021). The figure shows that the two lines were closely related to each other in the month before the treatment, and an immediate increase in the impactful count rate in the treatment group once the machine learning decision model for recommending counts was in place. On day one of the introduction of the new policy, the treatment group experienced a significant increase in impactful count rate, however, the ICR dropped on day two. We observed a similar trend on days 8, 15, 22. After closer examination, we realised these days corresponded to Sundays, suggesting that at least part of this drop is due to incorrect/incomplete counting at the stores due to employees experiencing lower morale and/or increased workload during Sundays.

Overall, our results show that the machine learning decision model provided a relative improvement of 26.07% in detecting inventory record inaccuracies with respect to those identified by operations analysts using heuristics and root cause analysis. This improvement means that stores in the treatment group detected and corrected inventory record errors swifter, helping stores to make more appropriate replenishment and assortment decisions. A closer look at the data, however, also reveals that the number of impactful errors detected by the machine learning decision model decreased over time.

This was particularly the case approximately 24 days after the field experiment started. While this relationship did not exist for recommendations done by analysts using traditional methods and heuristics, we speculate that the decay may be associated with the black Friday week, in which increased SKU demand may have affected the stores' labour availability run auditing routines. Further discussion on this phenomenon is presented in section 3.6.

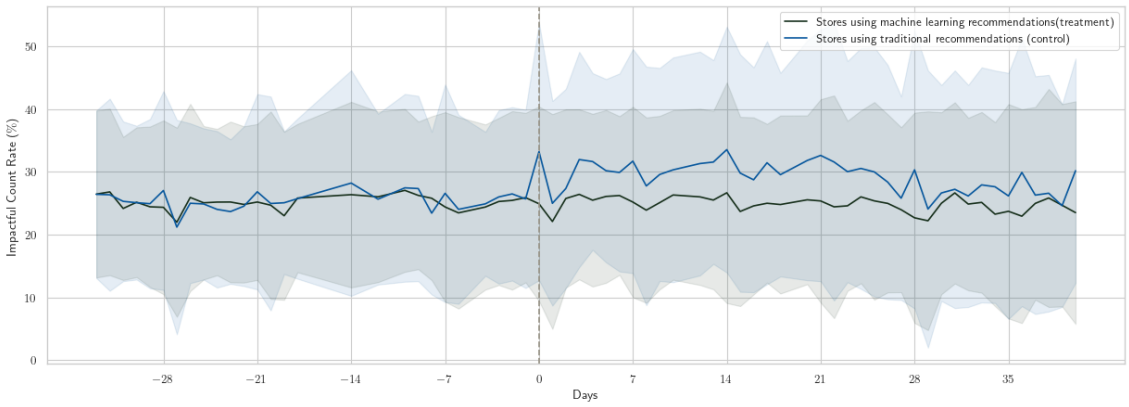


Figure 3-4 Number of impactful errors detected before and after the model was deployed for traditional and algorithm-recommended counts. The shaded area represents standard errors.

Because having more accurate inventory records helps stores make better replenishment decisions, we noted that stores in the treatment group had lower lost sales and waste. In Table 3-11, we formally measure the difference in performance on stores' inventory accuracy, sales, waste, and product availability using a deep learning versus traditional recommendation cycle count model. Our randomized experiment results confirm that stores in the treatment group are more likely to accurately detect inventory records in errors and so better inventory accuracy. We observed a 0.05% improvement in the sales forecast error and a 0.02% improvement in waste. While these numbers seem small, they improved the company's profitability in the order of millions. They also played a key role in improving the availability of the most important products.

Table 3-11 Overall effect of the deep learning-based model on performance

Impactful count rate	Pre-trial	Trial	Trial effect	Overall effect
Treatment group	25.1 %	31.4 %	6.2 %	6.7%*
Control group	25.7 %	25.4 %	-0.5%	
Waste cost as % of sales				
Treatment group	1.22 %	1.17%	-0.05%	-0.02%*
Control group	1.21%	1.18%	-0.03%	
Sales Forecast error				
Treatment group	14.50 %	15.41 %	0.92 %	-0.05%*
Control group	14.47%	15.44 %	0.97 %	
Sales Based Availability				
Treatment group	94.26%	94.00%	-0.26%	0.06%*
Control group	94.0%	93.68%	-0.32%	

*Note: * Significant at 0.1 level*

To facilitate the interpretation of our results, Figure 3-5 displays the net effect of our decision model on customer service level, calculated using as $1 - (Lost\ Sales / Adjusted\ Sales)$. The figures show that during the first week after the introduction of the machine learning decision model, there was no effect on customer service level, as any new orders corrected will take some time to arrive at the stores. After the first week $(t = 7)$, the effect of the deep-learning recommendations becomes salient. We noted that before our model was deployed, the customer service level reached negative points – customers were dissatisfied as some products were unavailable – after our decision model was implemented, the customer service level remained positive, and reached new maximum levels, which matches the sales gains experienced by Superstore.

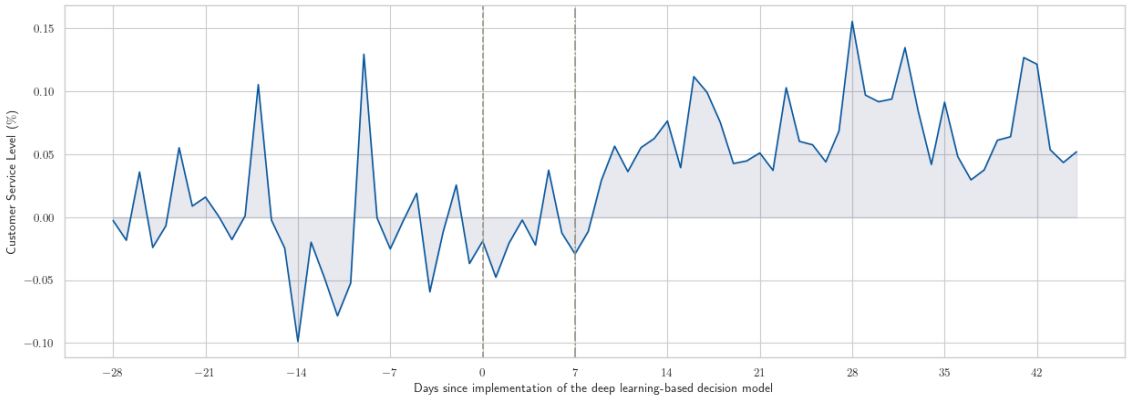


Figure 3-5 Net effect of the deep learning-based decision model on the customer service level.

3.5.2 Enactment of algorithmic recommendations

Interestingly, although algorithmic recommendations increase various operational performance indicators, they also lead to lower compliance in treated stores as shown in Figure 3-6. Workers in the control group performed on average 82% of the recommended counts, while stores in the treatment group only adhered on average to 78% of recommendations. The results from running a two-sample t-test highlight that the two distributions are moderately different with a value of 2.1073.

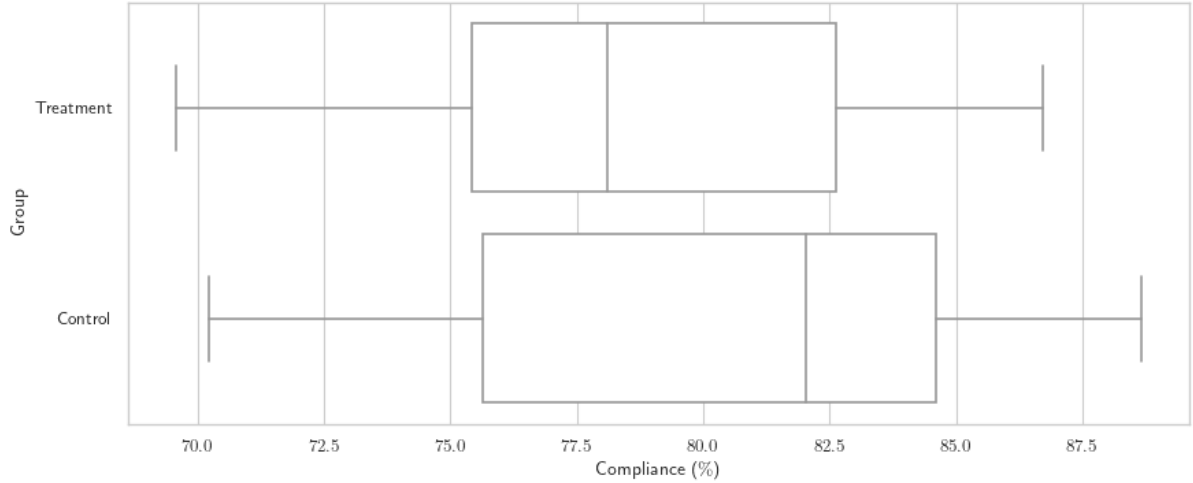


Figure 3-6 Comparison of counting recommendation adherence between control and treatment groups

3.5.3 Robustness of the effects

We performed several robustness checks related to our field experiment. This included comparing the model against the suspicious minimum heuristic, against other machine learning algorithms, and across stores.

Comparing predictions against suspicious minimum heuristic

Stores receiving count recommendations from analysts' heuristics exhibited a less reactive performance across the entire trial period, supporting the idea that these heuristics only represent a portion of the IRI problem. Consider, for example, the suspicious minimum heuristic that recommends stores counting a SKU if its records show a positive stock, but no sales, aiming to capture overstated inventory records when the on-hand inventory is zero. This heuristic recommends impactful counts almost all the time because the probability of having stock but not sales continuously would be extremely low (DeHoratius et al., 2008). Yet it only recognises one failure mode of IRI, which harms the overall detection of inventory records. While the company have created around 11 heuristics, and analysts also recommend counts based on them running root cause analysis, their recommendations still fail to cover as many patterns of error as the deep learning decision model recognised.

Comparing predictions against shallow machine learning algorithms

Because various machine learning algorithms may provide an alternative way to model IRI, we compared the deep learning results against seven machine learning classification algorithms, including logistic regression, random forests and AdaBoost classifiers. The comparison was done using the area under the receiving operating characteristic curve, or AUC for short (see Figure 3-7). Out of these algorithms, the Random Forest exhibited the best performance, yet it had a 13% lower performance than the deep learning-based decision model. Thus, we found that even random forest algorithms predicted more false

positives and false negatives than the deep learning method, supporting the notion that deep learning is more effective at representing the problem of IRI by distinguishing better between error categories.

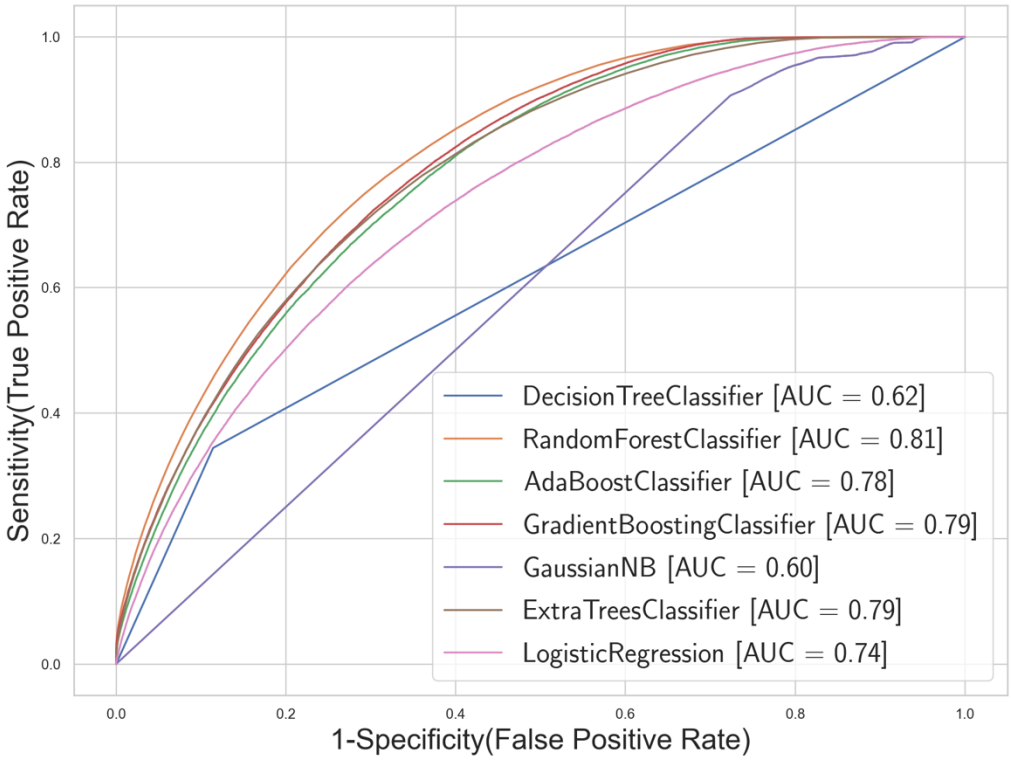


Figure 3-7 Receiver operating characteristic curve for various machine learning models considered in the study.

Comparing predictions across stores

We compared the performance between stores and the performance of different SKU categories. We first plotted a histogram of the percentage of impactful inventory errors detected across all the treatment stores during the trial period (see Figure 3-8). We noticed that while the model performance varied widely, it resembled a normal distribution centred around 31% ICR, suggesting the deep learning model generalises well across stores. We noted, however, that about 14% of the observations had an ICR lower

than 10%, suggesting some anomalous behaviour. Further analysis showed that those stores had a compliance rate lower than average; indicating that stores' employees ignored the recommended counts.

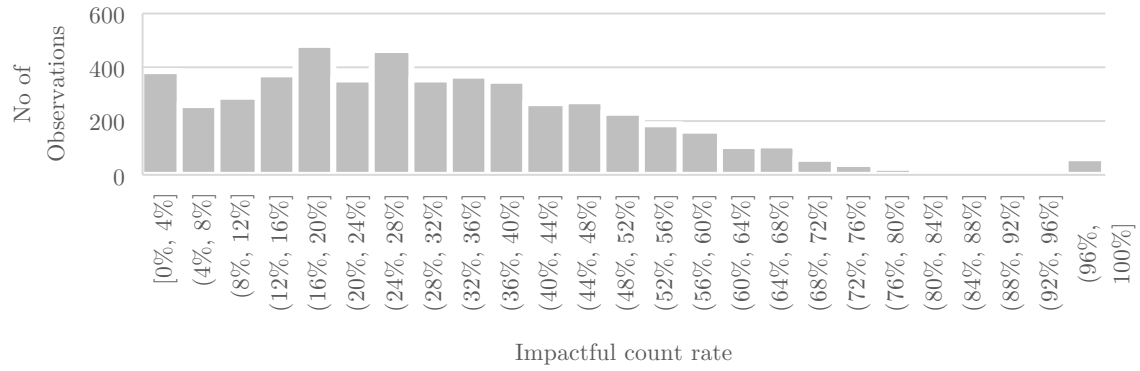


Figure 3-8 The percentage of impactful inventory errors detected across all treatment stores during the trial period.

3.6 Discussion

We have investigated the implementation of a deep learning algorithm to address the operational problem of inventory record inaccuracy (IRI) in a retail context. We show that a deep learning-based decision model can boost operational performance by providing better count recommendations to stores compared to traditional heuristics and policies for stock inspection and replenishment under imperfect information (e.g. Kök and Shang, 2007). Particularly, our experiment results show that stores that have deployed a deep learning decision model were able to detect discrepancies between recorded and in-hand stock 26.7% faster than those stores deploying traditional heuristics and analysts' recommendations. By improving inventory integrity, stores in the treatment group were able to make better decisions on when to replenish an SKU, which in turn led to measurable improvements in key sales and waste performance measures. In addition to enhancing operational performance, our findings are of theoretical interest as addressing IRI requires workers to enact algorithmic recommendations, which in turn

highlights the coupling of humans and algorithms in retail operations. In this section, we build on two diverse streams of literature: first, on the augmentation of operations with algorithmic technologies, and second, on the interplay between humans and algorithms. We build on these literatures to theorize the mechanisms that affect the realisation of improvements in work that require human actors to interface with algorithms. To that effect, we augmented our quantitative results with further qualitative research involving analysts, managers, and stock controllers and observations of stock controllers performing counts in supermarkets. In this section, we will discuss how our findings extend the existing literature toward defining three distinct mechanisms that mediate the interaction between humans and algorithms: representation, aversion, and degradation, which we will discuss in turn.

3.6.1 Representation

As reported above, we found that stores receiving algorithmic recommendations exhibited better operational performance than stores receiving recommendations from operations analysts' heuristics. The main reason is that the operation analyst heuristics fail to capture all the patterns of error that cause IRI; they capture some issues such as inventory "freezing" (DeHoratius et al., 2008), but they remain unable to identify obscure issues. The core challenge for operations analysts lies in the complexity inherent in identifying and correcting inventory record errors from large amounts of data. When the retailer implemented the deep learning algorithm, the identification of errors improved, for two reasons: For one, the algorithm represents the IRI problem in multiple layers, synthesising patterns of errors that the operation analyst could not find. Our results here are consistent with previous literature (e.g. Murray et al., 2021) which posits that using algorithmic technologies would increase the organisation's information processing capacity and minimise the shortcomings of human decision-making, including 'myopic

search' and 'bounded rationality' (Gavetti et al., 2012; March et al., 1958; Simon, 1947; Simon, 1991). Our results support Gavetti and Levinthal's (2000) general notion that the ability to create better problem representations enables actors to process more information and focus their attention on some aspects of the environment, ultimately driving more effective decision-making and performance.

A closer look at the human-algorithm ensembles in our study, however, reveals a somewhat more complex dynamic between representations and performance than that suggested by Gavetti and Levinthal. Specifically, we found that while algorithmic technologies may be able to create more complex problem representations, organisations can fail to convert these representations into better operational improvement. In our case, the deep learning representations existed as a complex set of weights and variables that were interconnected and continuously updated by an algorithm. These representations led to the identification of more data patterns of impactful errors and produced better results than operations analysts' heuristics. If these better representations were the sole factor in improving performance, we would have expected that stores using our deep learning decision model achieved a similar rate of impactful error identification as we saw in the model validation. Instead, we found that stores using the deep learning model only realised 46.37% of the expected improvement. Thus, representations are unlikely to be the main ingredient in the process of learning, as suggested by Gavetti and Levinthal (2000). This is particularly the case when decision and enactment are divided across algorithms and human actors. Our findings demonstrate that decision-making is distributed, not only between actors but also between these actors and material technologies capable of exhibiting intentionality.

3.6.2 Aversion

A second mechanism that influences the effectiveness of recommender systems is algorithmic aversion (Dietvorst et al., 2015). This term refers to the dislike or avoidance of algorithmic recommendations and has conventionally been linked to human reactions when algorithms make errors (Dietvorst et al., 2015, 2018). Put simply, people tend to lose trust in algorithms if they make mistakes, even if the algorithm's recommendations are on average superior to human ones. More recently, researchers have contended that a second type of algorithmic aversion exists, one that transcends reliance on recommendation accuracy. This form of aversion – that we call affective - arises because humans generally prefer recommendations from other humans rather than relying on algorithms (Yeomans et al., 2019).

Considering these studies, one might assume that algorithmic aversion would not be present in situations where individuals are unaware of whether recommendations come from a recommender system or a human, especially when the algorithmic recommendations are more accurate. However, our findings present a contrasting narrative. We found that individuals, without their awareness, exhibit a preference for less accurate human recommendations over more precise algorithmic ones.

This is surprising as we designed our experiment to uphold participant blinding in both the treatment and control cohorts regarding the instantiation of the deep learning decision model and their subsequent group allocation. Careful measures were implemented to preclude algorithmic aversion, which might arise due to emotional reactions such as anxiety, fear, or mistrust towards algorithmic recommendations, thereby safeguarding the integrity of the findings. Astonishingly, despite the discernible superiority of the model's recommendations over those proffered by the operation

analysts (as demonstrated in section 3.4.2.2), the stock controllers demonstrated a notable inclination towards the less precise human recommendations.

This finding leads us to propose the existence of a distinct form of algorithmic aversion—one that hinges upon the specific content of the recommendations being offered. Our analysis suggests that this form of aversion arose from stock controllers' adaptation to the new counts' recommendations. The deep learning algorithm's recommended counts were more challenging to perform due to their inclusion of SKUs with higher stock levels, frozen items, or visually similar adjacent SKUs. As a result, stock controllers chose to skip some of these counts as they required more time or caused personal discomfort.

To supplement and corroborate our findings regarding aversion, we employed an observational approach, shadowing stock controllers during their daily counts. Our observations, albeit limited, supported this conclusion. We found that stock controllers assessed whether they could perform a count within the allocated time frame and if the SKUs were easily accessible. If any of these conditions were not met, they often opted to skip the count rather than perform it poorly, expressing their desire not to "mess anymore with the stock control system." (quote from stock controller #1).

To gain greater insights into the distinctions between the proposed aversion and other potential behaviours, such as quota restrictions and goldbricking, we draw upon Roy's (1952) seminal study in industrial relations, which provides valuable insights into the examination of stock controllers at Superstore. Roy's research on workers in a machine shop elucidates their rational behaviour, governed by two norms: quota restrictions and goldbricking. Quota restrictions involve individuals restraining themselves from meeting predetermined quotas due to concerns that meeting or exceeding them might lead to inflated expectations from management or reduced allocated time (or

resources) for specific tasks. On the other hand, goldbricking entails individuals intentionally exerting less effort than their capabilities while maintaining the illusion of productivity, particularly when confronted with more arduous or monotonous tasks.

Although my observations of stock controllers adhering to the targeted count routine were limited, they suggest indications of both quota restrictions and goldbricking. Throughout my observations, I found that achieving a 100% counting rate was indeed possible (see Figure 3-9 below), as supported by the statistical analysis of non-adherence to recommendations. However, the data also disclosed that in over 50% of cases, stock controllers did not fully comply with the recommendations. As a result, we propose that while stock controllers could adhere to all recommendations, it is highly likely that they engage in some form of quota restrictions, intentionally limiting their efforts to stay below the established target. By consistently adhering to 100% of the recommendations, management might revise and curtail the time allocated for completing the routine.

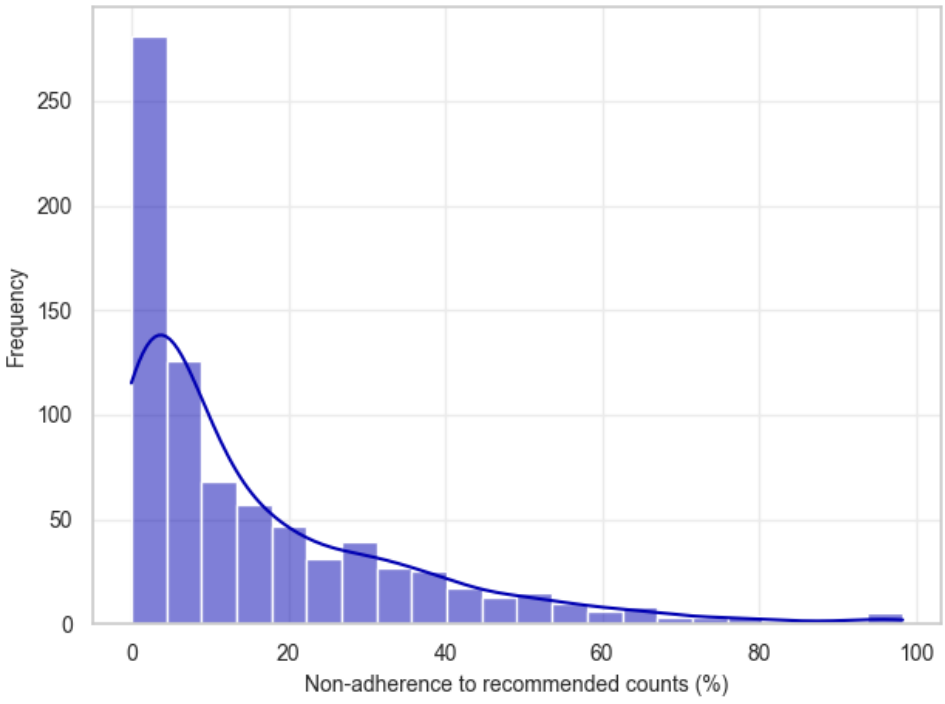


Figure 3-9 Store's non-adherence to counting recommendations.

Furthermore, our observations revealed that stock controllers engaged in goldbricking behaviour, particularly when the recommendations involved more difficult and tedious counting tasks. This behaviour aligns with Roy's (1952) study on goldbricking, where machine shop workers avoided "stinker" jobs that required additional effort. In our study at Superstore, we discovered a similar pattern. Stock controllers limited their performance when the recommendations involved products that were particularly challenging to count, such as frozen items, liqueurs, and products with high inventory levels.

For stock controllers, performing a difficult count offers no additional benefits compared to an easy count. They face penalties for undertaking such tasks. They must endure chilling temperatures in refrigerated areas (often avoiding the use of protective gloves as they hinder their ability to manoeuvre the stock cages), and they become frustrated when they are unable to accurately count all the products. Additionally, some products were deliberately left uncounted due to store managers discouraging employees from performing counts that could result in reporting waste. Waste serves as an important key performance indicator for the store, and Superstore management utilizes it to allocate budgets. Consequently, managers reinforce cultural norms that discourage reporting waste as it could impact the budget for the following year, which potentially prevents them from counting products with high inventory levels.

We argue that these penalties and discouragement condition the behaviour of stock controllers, leading them to avoid counting certain difficult-to-count products. This creates a recursive cycle between aversion and goldbricking. Consequently, a norm of goldbricking emerges, whereby stock controllers tend to slack off on difficult-to-count products. This norm becomes socialized, influencing even new stock controllers to learn

and adopt it, which, in turn, makes them averse to certain types of recommendations, regardless of the source providing them.

3.6.3 Degradation

The third mechanism affecting the realisation of operational improvements was degradation. By 'degradation' we mean the process through which the quality of the algorithm recommendations decays over time. Our use of the term degradation is, thus, analogous to the concept of model drift in data science. We use degradation in this context to emphasize the perceived quality of the algorithm's recommendations, in the eyes of the person interfacing with the algorithm in question. Degradation poses a problem because it leads the algorithm to create representations of the problem that are increasingly less reflective of the real-world environment. For example, in our case, we noted that as time passed, the deep learning algorithm tended to recommend less useful counts. This drop in quality may be caused by changes in seasonality patterns that lead to new data distributions, and in turn, less useful recommendations. Degradation is usually overcome by retraining the algorithm continuously on data that reflect new changes in the environment. This is more difficult in practice, however, as it may take some time for changes in the environment to be significantly represented in the data used for retraining the model. Hence, there will inevitably be periods in which the algorithm recommendations are offset with the ground truth, increasing the risk exposure for organisations relying on these systems to make decisions.

From shadowing stock controllers, it was also apparent that the problem of degradation was more significant than our study findings suggested as stock controllers often pretend to adhere to the algorithm recommendations, ultimately affecting the quality of the data used to retrain these models. The evidence for this was the observation that some stock controllers avoid reporting the lack of adherence to recommended counts.

Instead, they engaged in a practice that they call 'screen counting' through which they mislead the system to report they performed the count by simply entering (i.e., falsely confirming) the last stock quantity registered in the system This deceptive behaviour severely impacted the data quality used for retraining the models, consequently affecting the accuracy of their predictions. The introduction of uncertainty into what we believed to be ground truth labels further aggravated the degradation of the solution (Lebovitz et al., 2021).

Even after deploying our system to all the UK stores and subjecting it to continuous weekly retraining, conversations revealed that degradation persisted as a critical problem for the company. This observation indicates a clear link between the degradation phenomenon and aversion. As such, both factors need to be comprehensively understood and addressed in tandem to effectively combat the issue.

In human-algorithm ensembles, the identification of degradation in general will be challenging, in particular in cases where the human receiving the recommendation has low discretion. While professionals with high discretion may use their expertise and status to challenge the algorithm's outputs (Anthony, 2018, 2021; Lebovitz et al., 2022), this level of questioning would be prohibitive for employees with low levels of discretion, such as stock controllers. In these cases, the identification of degradation requires monitoring practices through which the performance of the human algorithm ensemble is tracked over time.

3.6.4 Managerial implications

The rising availability of detailed and timely operational data and the growing openness among companies to explore digital technologies are driving the adoption of general-purpose algorithmic technologies to improve operations (e.g. Bastani et al., 2020; Mišić and Perakis, 2020; Senoner et al., 2022). Our research contributes a new method to a

growing portfolio of studies that demonstrate the potential for deep learning algorithms for operational planning and control tasks. Our decision model can be efficiently adopted for inventory control problems, where it has shown superior performance compared to traditional inventory correction strategies. We showed that by detecting and correcting IRI faster, the deep learning model also helped managers make better planning and control decisions, leading to improvements in key performance indicators such as sales and waste. In the case at hand, the estimated increase in sales has been estimated by Superstore as tens of millions per annum.

Our field experiment was conducted in a grocery and merchandise retailer, the robustness checks indicated that our model could be applied to address IRI for different stock-keeping unit categories, providing strong evidence that our model can be successfully transferred to other retail and manufacturing contexts more generally.

A second practical implication is based on our observation that generic machine learning algorithms can outperform even experienced analysts when addressing operational problems. We noticed, for example, that companies benefit most from using algorithmic decision models for poorly understood problems, or those that depend on many variables, say IRI, even if the company currently employ workers with high domain expertise in their respective field, say production planning and control. At the same, algorithmic solutions have important limitations. While these models may improve performance, they can also degrade over time and suffer from algorithmic 'gaming'. The implementation of algorithmic technologies must thus be accompanied by human oversight in the development, curation and monitoring of any algorithmic solutions employed for operational problems. Data science teams should work closely with domain experts to define patterns and solutions to the problem of users taking strategic actions to game the algorithmic recommendations. In the case of IRI, for example, retailers can

align the incentives of the store employees to improve compliance with the deep learning recommendations. A further strategy to increase the accuracy of the predictions could be to create ensemble recommendations from multiple meta-models to the point that some inventory errors could even be automatically corrected without inspecting the on-hand inventory. Such an approach would both decrease the workload for staff in the stores and increase compliance even for tedious, difficult or uncomfortable count recommendations.

3.7 Conclusion

Our study contributes to a growing body of knowledge that demonstrates the potential of algorithmic technologies to improve operational decision-making. Previous studies have demonstrated the potential and challenges of algorithmic recommendations in high-discretion settings (Allen and Choudhury, 2021; Burton et al., 2020; Dietvorst et al., 2018; Khosrowabadi et al., 2022; Lebovitz et al., 2022). We add a further aspect to this body of knowledge by demonstrating that algorithmic technologies can provide improvements even in low discretion settings that require human-algorithm collaboration. In the retail context we studied, this was demonstrated by the fact that stores receiving deep learning recommendations on which SKUs to count featured better inventory record accuracy, increased sales, and reduced waste. We argue this improvement stems partially from the fact that algorithms can create better representations of problems, thereby overcoming known human shortcomings such as bounded rationality and myopic search. However, our results also highlight the fact that better representations do not automatically translate into better performance in settings where the loci of the representation and behaviour differ. Here algorithmic aversion resulting from human compliance as well as algorithm degradation due to errors creeping into the model play a key role in moderating operational performance. This suggests that

in the short term, algorithmic technologies produce outstanding improvement that surpasses cumulative capabilities and yet, in the long run, the performance may backslide if the models are not attended to, monitored and updated. More broadly, we thus show how generic deep learning algorithms can be successfully applied to operations management problems, provided one is cognisant of their inherent limitations and the potential human behaviours that can ensue on the back of their implementation.

References

Allen, R. & Choudhury, P. (2021). Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion. *Organization Science, 33* (1), 149-169.

Anthony, C. (2018). To question or accept? How status differences influence responses to new epistemic technologies in knowledge work. *Academy of Management Review, 43* (4), 661-679.

Anthony, C. (2021). When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative Science Quarterly, 66* (4), 1173-1212.

Bastani, H., Zhang, D. & Zhang, H. (2020). Applied machine learning in operations management. *Available at SSRN 3736466,*

Bernard, P. (1985). Cycle counting. *Production Inventory Management, 26* (4), 27-41.

Bernstein, F., Modaresi, S. & Sauré, D. (2019). A dynamic clustering approach to data-driven assortment personalization. *Management Science, 65* (5), 2095-2115.

Brooks, R.B. & Wilson, L.W. (2007). *Inventory record accuracy: unleashing the power of cycle counting*, 2nd ed. John Wiley & Sons, New York.

Brynjolfsson, E. & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review, 98* (4), 1-20.

Brynjolfsson, E. & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science, 358* (6370), 1530-1534.

Burton, J.W., Stein, M.K. & Jensen, T.B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making, 33* (2), 220-239.

Chawla, N.V., Bowyer, K.W., Hall, L.O. & Kegelmeyer, W.P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research, 16*, 321-357.

Chollet, F. (2018). *Deep learning with Python*, 1st ed. Manning Publications, Shelter Island.

Chou, Y.C., Chuang, H.H.C., Chou, P. & Oliva, R. (2022). Supervised Machine Learning for Theory Building and Testing: Opportunities in Operations Management. *Journal of Operations Management, 69* (4), 643-675.

Choudhury, P., Allen, R.T. & Endres, M.G. (2021). Machine learning for pattern discovery in management research. *Strategic Management Journal, 42* (1), 30-57.

Chuang, H.H.-C. (2015). Mathematical modeling and Bayesian estimation for error-prone retail shelf audits. *Decision Support Systems, 80*, 72-82.

Chuang, H.H.-C. & Oliva, R. (2015). Inventory record inaccuracy: Causes and labor effects. *Journal of Operations Management, 39* (1), 63-78.

Chuang, H.H.-C., Oliva, R. & Liu, S. (2016). On-shelf availability, retail performance, and external audits: A field experiment. *Production and Operations Management, 25* (5), 935-951.

Chuang, H.H.C., Oliva, R. & Kumar, S. (2022). Group-level information decay and inventory inspection: An empirical-analytical approach. *Journal of Operations Management, 68* (2), 130-152.

Corsten, D. & Gruen, T. (2004). Stock-outs cause walkouts. *Harvard Business Review, 82* (5), 26-28.

Csaszar, F.A. & Steinberger, T. (2022). Organizations as Artificial Intelligences: The Use of Artificial Intelligence Analogies in Organization Theory. *Academy of Management Annals, 16* (1), 1-37.

Curchod, C., Patriotta, G., Cohen, L. & Neysen, N. (2020). Working for an algorithm: Power asymmetries and agency in online work settings. *Administrative Science Quarterly, 65* (3), 644-676.

Cyert, R.M. & March, J.G. (1963). *A behavioral theory of the firm.* Prentice-Hall, Englewood Cliffs.

De Baets, S. & Harvey, N. (2018). Forecasting from time series subject to sporadic perturbations: Effectiveness of different types of forecasting support. *International Journal of Forecasting, 34* (2), 163-180.

DeHoratius, N., Mersereau, A.J. & Schrage, L. (2008). Retail inventory management when records are inaccurate. *Manufacturing & Service Operations Management, 10* (2), 257-277.

DeHoratius, N. & Raman, A. (2008). Inventory record inaccuracy: An empirical analysis. *Management Science, 54* (4), 627-641.

Delaunay, C., Sahin, E. & Dallery, Y. (2007). *A literature review on investigations dealing with inventory management with data inaccuracies*, 2007 1st Annual RFID Eurasia. IEEE, pp. 1-7.

Dietvorst, B.J., Simmons, J.P. & Massey, C. (2015). Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144* (1), 114.

Dietvorst, B.J., Simmons, J.P. & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science, 64* (3), 1155-1170.

Emma, C.K. (1966). Physical Inventory Procedures at Navy Stock Points-Observations on Physical Inventory and Stock Record Accuracy. *Naval Supply Research and Development Facility, Bayonne,*

Eroglu, C. & Croxton, K.L. (2010). Biases in judgmental adjustments of statistical forecasts: The role of individual differences. *International Journal of Forecasting, 26* (1), 116-133.

Faraj, S., Pachidi, S. & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization, 28* (1), 62-70.

Ferreira, K.J., Lee, B.H.A. & Simchi-Levi, D. (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management, 18* (1), 69-88.

Gardner, P.H. & Berry, D.C. (1995). The effect of different forms of advice on the control of a simulated complex system. *Applied Cognitive Psychology, 9* (7), S55-S79.

Gavetti, G., Greve, H.R., Levinthal, D.A. & Ocasio, W. (2012). The behavioral theory of the firm: Assessment and prospects. *Academy of Management Annals, 6* (1), 1-40.

Gavetti, G. & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly, 45* (1), 113-137.

Glaser, V.L., Pollock, N. & D'Adderio, L. (2021). The biography of an algorithm: Performing algorithmic technologies in organizations. *Organization Theory, 2* (2), 26317877211004609.

Hardgrave, B.C., Aloysius, J.A. & Goyal, S. (2013). RFID-enabled visibility and retail inventory record inaccuracy: Experiments in the field. *Production and Operations Management, 22* (4), 843-856.

Hardgrave, B.C., Goyal, S. & Aloysius, J.A. (2011). Improving inventory management in the retail store: The effectiveness of RFID tagging across product categories. *Operations Management Research, 4* (1-2), 6-13.

Heese, H.S. (2007). Inventory record inaccuracy, double marginalization, and RFID adoption. *Production and Operations Management, 16* (5), 542-553.

Holland, C., Levis, J., Nuggehalli, R., Santilli, B. & Winters, J. (2017). UPS optimizes delivery routes. *Interfaces, 47* (1), 8-23.

Holweg, M., Disney, S., Holmström, J. & Småros, J. (2005). Supply Chain Collaboration. *European Management Journal, 23* (2), 170-181.

Iglehart, D.L. & Morey, R.C. (1972). Inventory systems with imperfect asset information. *Management Science, 18* (8), B-388.

Imdahl, C., Hoberg, K. & Schmidt, W. (2021). Targeted automation of order decisions using machine learning. *Available at SSRN 3822131,*

Kang, Y. & Gershwin, S.B. (2005). Information inaccuracy in inventory systems: stock loss and stockout. *IIE transactions, 37* (9), 843-859.

Kellogg, K.C., Valentine, M.A. & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals, 14* (1), 366-410.

Khosrowabadi, N., Hoberg, K. & Imdahl, C. (2022). Evaluating human behaviour in response to AI recommendations for judgemental forecasting. *European Journal of Operational Research, 303* (3), 1151-1167.

Kök, A.G. & Shang, K.H. (2007). Inspection and replenishment policies for systems with inventory record inaccuracy. *Manufacturing & Service Operations Management, 9* (2), 185-205.

Lebovitz, S., Levina, N. & Lifshitz-Assaf, H. (2021). Is AI ground truth really "true"? The dangers of training and evaluating AI tools based on experts' know-what. *Management Information Systems Quarterly, 45* (3), 1501-1526.

Lebovitz, S., Lifshitz-Assaf, H. & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science, 33* (1), 126-148.

LeCun, Y., Bengio, Y. & Hinton, G. (2015). Deep learning. *Nature, 521* (7553), 436-444.

Lee, H. & Özer, Ö. (2007). Unlocking the value of RFID. *Production and Operations Management, 16* (1), 40-64.

Lee, H.L., Padmanabhan, V. & Whang, S. (1997). The bullwhip effect in supply chains. *Sloan Management Review, 38* (3), 93-102.

Leonardi, P.M. & Barley, S.R. (2008). Materiality and change: Challenges to building better theory about technology and organizing. *Information and Organization, 18* (3), 159-176.

Levy, M., Weitz, B.A., Grewal, D. & Madore, M. (2004). *Retailing management.* McGraw-Hill Irwin, New York.

March, J.G., Simon, H.A. & Guetzkow, H.S. (1958). *Organizations.* Wiley, New York.

Mišić, V.V. & Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management, 22* (1), 158-169.

Morales, D.R. & Wang, J. (2010). Forecasting cancellation rates for services booking revenue management using data mining. *European Journal of Operational Research, 202* (2), 554-562.

Morey, R.C. & Dittman, D.A. (1986). Optimal timing of account audits in internal control. *Management Science, 32* (3), 272-282.

Murray, A., Rhymer, J. & Sirmon, D.G. (2021). Humans and technology: Forms of conjoined agency in organizations. *Academy of Management Review, 46* (3), 552-571.

Oliva, R. & Watson, N. (2009). Managing functional biases in organizational forecasts: A case study of consensus forecasting in supply chain planning. *Production and Operations Management, 18* (2), 138-151.

Orlikowski, W.J. & Scott, S.V. (2008). 10 sociomateriality: challenging the separation of technology, work and organization. *Academy of Management Annals, 2* (1), 433-474.

Orlikowski, W.J. & Scott, S.V. (2015). The algorithm and the crowd. *MIS Quarterly, 39* (1), 201-216.

Pachidi, S., Berends, H., Faraj, S. & Huysman, M. (2021). Make way for the algorithms: Symbolic actions and change in a regime of knowing. *Organization Science, 32* (1), 18-41.

Porac, J.F., Thomas, H. & Baden-Fuller, C. (1989). Competitive groups as cognitive communities: The case of Scottish knitwear manufacturers. *Journal of Management Studies, 26* (4), 397-416.

Raisch, S. & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review, 46* (1), 192-210.

Raman, A., DeHoratius, N. & Ton, Z. (2001). Execution: The missing link in retail operations. *California Management Review, 43* (3), 136-152.

Redman, R. (2021). *NielsenIQ: Out-of-stocks during pandemic lead to billions in missed CPG sales.* Available at: https://www.supermarketnews.com/consumer-trends/nielseniq-out-stocks-during-pandemic-lead-billions-missed-cpg-sales (Accessed: April 5th 2021)

Rinehart, R.F. (1960). Effects and causes of discrepancies in supply operations. *Operations Research, 8* (4), 543-564.

Roy, D. (1952). Quota restriction and goldbricking in a machine shop. *American Journal of Sociology, 57* (5), 427-442.

Sarac, A., Absi, N. & Dauzère-Pérès, S. (2010). A literature review on the impact of RFID technologies on supply chain management. *International Journal of Production Economics, 128* (1), 77-95.

Schrady, D. (1970). Operational definitions of inventory record accuracy. *Naval Research Logistics Quarterly, 17* (1), 133-142.

Scott, S.V. & Orlikowski, W.J. (2014). Entanglements in practice. *MIS Quarterly, 38* (3), 873-894.

Senoner, J., Netland, T. & Feuerriegel, S. (2022). Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science, 68* (8), 5704-5723.

Sheppard, G.M. & Brown, K.A. (1993). Predicting inventory record-keeping errors with discriminant analysis: A field experiment. *International Journal of Production Economics, 32* (1), 39-51.

Shrestha, Y.R., Ben-Menahem, S.M. & Von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review, 61* (4), 66-83.

Simon, H.A. (1947). *Administrative behavior : a study of decision-making processes in administrative organization*. Macmillan, New York.

Simon, H.A. (1991). Bounded rationality and organizational learning. *Organization Science, 2* (1), 125-134.

Sun, J., Zhang, D.J., Hu, H. & Van Mieghem, J.A. (2021). Predicting human discretion to adjust algorithmic prescription: A large-scale field experiment in warehouse operations. *Management Science,*

Thagard, P. (2005). *Mind: Introduction to cognitive science*. MIT press,

Ton, Z. & Raman, A. (2010). The effect of product variety and inventory levels on retail store sales: A longitudinal study. *Production and Operations Management, 19* (5), 546-560.

Uçkun, C., Karaesmen, F. & Savaş, S. (2008). Investment in improved inventory accuracy in a decentralized supply chain. *International Journal of Production Economics, 113* (2), 546-566.

van der Haar, J., Wellens, A.P., Boute, R.N. & Basten, R.J. (2023). Supervised Learning for Integrated Forecasting and Inventory Control. *Available at SSRN,*

Van Donselaar, K.H., Gaur, V., Van Woensel, T., Broekmeulen, R.A. & Fransoo, J.C. (2010). Ordering behavior in retail stores and implications for automated replenishment. *Management Science, 56* (5), 766-784.

Waardenburg, L., Huysman, M. & Sergeeva, A.V. (2021). In the land of the blind, the one-eyed man is king: Knowledge brokerage in the age of learning algorithms. *Organization Science, 33* (1), 59-82.

Weick, K.E. (1979). *The social psychology of organizing*, 2d ed. McGraw-Hill, Inc., New York.

Yeomans, M., Shah, A., Mullainathan, S. & Kleinberg, J. (2019). Making sense of recommendations. *Journal of Behavioral Decision Making, 32* (4), 403-414.

Chapter 4

Using algorithmic technologies to improve knowledge work

Abstract

In this study, we explore how organisations leverage algorithms to improve knowledge work in tasks that require skilled performances. Drawing on a multiple-case study of four business areas in a multinational energy firm undergoing digital transformation, we found two distinct pathways for transforming knowledge work. The first pathway, “altering a specific task”, involved a one-on-one substitution of humans with algorithms. When following this pathway in tasks requiring skilled performances, business areas experienced complications that affected both task execution and human learning from the algorithm outputs, as the literature would predict. Nonetheless, we also found a second pathway, “re-engineering an entire process”, wherein adjacent sequences of steps to the task are redesigned upon algorithm integration. This pathway allowed business areas to create improvements by deploying algorithms in tasks that demanded skilled performances. We also found that successful algorithm adoption requires an "alteration of the knowledge regime" to adapt evaluation schemes and ensure the quality of algorithm outputs. Building upon these findings, we propose a process model for the adoption of algorithmic technologies in knowledge work. Contrary to prevailing discourse, our research challenges the notion that the need for creativity, social intelligence and physical manipulation skills are the primary determinants of improving knowledge work with algorithmic technologies. Instead, we argue that this is a function of the organisations' ability to capture, translate, and modify the underlying know-how of the practice.

***Keywords*:** Knowledge work, process improvement, automation, machine learning, robotic process automation.

4.1 Introduction

The integration of new algorithmic technologies in knowledge work, a form of work reliant on the mastery of information, expertise and cognitive skills for value creation (Bechky, 2006; Davenport, 2015; Starbuck, 1992), assumes a pivotal role in the context of the digital transformation, with profound consequences on operational performance. These technologies encompass computer-programmed procedures engineered to emulate human work and decision-making (Kellogg et al., 2020), with machine learning being a notable example where algorithms learn to make predictions by discerning patterns in existing data (Agrawal et al., 2018). In recent years, organisations have sought to embrace new algorithms to restructure knowledge work across professional, clerical and managerial services (Armour and Sako, 2020; Sampson, 2021). However, the extent to which algorithms effectively drive operational performance in knowledge work remains a point of contention (Brynjolfsson et al., 2019). A fundamental challenge emerges when replacing knowledge workers with algorithms, particularly in tasks that necessitate the demonstration of knowledge; the '*ability to draw distinctions, within a collective domain of action, based on an appreciation of context or theory or both*' (Tsoukas, 2005, p.123). Such tasks include exercising specialized judgment and tailoring solutions to address custom problems and often serve as the source of value creation in organisations (Harvey, 2016; Staats et al., 2011). Thus, replacing knowledge work with algorithms may invertedly impede organisational learning and operational performance (Balasubramanian et al., 2022).

The available research on algorithmic technologies provides a partial understanding of their effects on improving knowledge work. Despite the widespread recent interest in algorithmic technologies under the labels of '*learning algorithms*' (Faraj et al., 2018), '*modern epistemic technologies*' (Anthony, 2018), '*intelligent technologies*' (Bailey and Barley, 2020), and '*artificial intelligence*' (Lebovitz et al., 2021), we still do not know how organisations can harness algorithmic technologies to enhance knowledge work and the specific circumstances under which this is possible. Prevailing economic studies often

assume that new technologies primarily contribute to productivity growth in routine work (e.g. Autor and Salomons, 2018; Autor et al., 2003), largely overlooking the potential impact of algorithms on non-routine knowledge work, albeit with a few exceptions (e.g. Frey and Osborne, 2017). In contrast, sociological and learning studies on algorithmic technologies, while contemplating such possibilities, do focus on the implications (and perils) of using algorithmic technologies for the construction of knowledge (Anthony, 2018; Lebovitz et al., 2021; Waardenburg et al., 2021), thereby leaving out the possibility that algorithms can play an active role in improving knowledge work. The question that remains, and is the focus of our investigation, is how organisations use algorithmic technologies to instantiate the improvement of knowledge work that demands skilled performances – akin to how process improvement is traditionally defined in operations management.

To answer this question, we conduct a multiple-case study design to explore and analyse the adoption of algorithms aimed at improving knowledge work in four business areas within a multinational firm, hereafter called *ACME*. We selected these areas to exemplify Davenport's (2015) four segments of knowledge work: *Business services* (as 'transactional work'), *internal audit* (as 'collaborative work'), *maintenance* (as 'integrative work') and *drilling* (as 'expert work'). By encompassing a diverse range of knowledge work, our study observed the influence of work variations on algorithmic utilization (Bailey and Barley, 2020). While our unit of observation was these business areas, we directed our analysis towards the foundational unit of work within these areas - the tasks performed. These tasks are tangible, dynamic, and context-dependent, constituting an essential part of the daily routines of the individuals involved. Accordingly, we built rich cases that capture knowledge work at the 'coalface' (in Barley (2008) parlance), highlighting the interplay between algorithmic technologies and the work contexts in which they are deployed (DeSanctis and Poole, 1994; Faraj and Pachidi, 2021; Orlikowski and Barley, 2001). To achieve this, we conducted structured interviews with members and managers in these areas, supplemented by participant observations

and analysis of internal performance reports. Within each case, we examined in greater detail one or two important efforts to transform tasks by adopting projects and constructed what we refer to as 'mini cases' or 'projects' for simplicity. Through these projects, we attended to how variations in work influenced the co-construction of knowledge between humans and algorithms, and the resulting impact on performance outcomes.

Overall, our findings make three contributions to the nascent debate on algorithmic technologies in knowledge-based work: first, we find that knowledge workers engage in a dynamic process to transform their work that, in turn, enables the full potential of algorithmic technologies. Our findings support that the knowledge workers studied enacted cognitive (Tripsas and Gavetti, 2000), and internal dynamics (Truelove and Kellogg, 2016), which shaped the adaptation to technological change, as well as an adjustment of the knowledge regime (Pachidi et al., 2021). This allowed them to increase the efficiency of knowledge work processes, while at the same time keeping control of the quality of the constructed knowledge and learning. This contribution is significant, as it contrasts with previous studies that have suggested that the opacity[6] of algorithmic recommendations (Lebovitz et al., 2022; Waardenburg et al., 2021), its embedded epistemic assumptions (Anthony, 2018; Lebovitz et al., 2021) and the stickiness of knowledge in social relations (Balasubramanian et al., 2022) prevent organisations from realising the potential of algorithmic technologies.

Second, we provide a differentiated view of the existing work on the automation of knowledge work (Acemoglu and Restrepo, 2020; Frey and Osborne, 2017; Sampson, 2021) by unpacking variations of work and technology implementation so far obscured by economic studies of work abstractions in databases (Bailey and Barley, 2020). Rooting our findings in the experience of four knowledge work organisations using algorithmic technologies, we highlight that the ability to use algorithms in knowledge work depends on the ability to capture, translate and modify the know-how of the practice, rather than

[6] Various algorithmic technologies exhibit varying degrees of opacity. However, it is widely acknowledged that all of them possess some level of opacity to the end-user, who typically lacks expertise in programming.

on the skills required to perform it. We draw on the notion of knowing-in-practice (Nicolini, 2011; Orlikowski, 2002) to explain that tasks requiring the same skill, may call for different boundaries to be bridged, and to lay out a set of propositions that explain mechanisms used to bridge these boundaries.

Finally, we combine our insights into how emergent algorithmic technologies alter work content and processes (Bailey et al., 2019) by proposing a conceptual model that depicts how organisations use algorithmic technologies to construct improvement opportunities. Overall, our findings contribute to the wider debate on the digital transformation of work, informing scholars and practitioners alike, towards a better understanding and predicting the impact that algorithmic technologies can have on the various kinds of knowledge work.

4.2 Theoretical background

4.2.1 Algorithmic technologies in improving knowledge work

Since the second half of the 20th century, knowledge work and technology have risen to become central themes within post-bureaucratic organisations (Barley and Kunda, 2001). Unlike manufacturing work, those who do 'knowledge' work engage in mental, rather than physical labour (Bechky, 2006); they manipulate knowledge and information to create value for the firm (Davenport, 2015; Starbuck, 1992). Formally defined, knowledge work involves the demonstration of knowledge; the "ability to draw distinctions, within a collective domain of action, based on an appreciation of context or theory or both" (Tsoukas, 2005). This type of work[7] plays a pivotal role in various occupations, including

[7] Although the terms "work" and "processes" are closely related, we acknowledge some key differences between them. "Work" primarily emphasizes the concrete, dynamic, and situated activities that an individual engages in as part of a living (Barley, 1996; Barley and Kunda, 2001; Orr, 1996). On the other hand, a "process" refers to a sequence of steps designed to transform inputs into economic outputs (Holweg et al., 2018). While there are instances where an individual's work may encompass an entire process, in modern organisations, processes are often divided among different individuals. Consequently, it becomes essential to differentiate between the concepts of "work" and "process".

healthcare professionals (Harvey, 2016), software developers (Staats et al., 2011) and financial analysts (Knorr Cetina, 2010). In these contexts, knowledge is *'(…) the critical input in production and primary source of value'* (Grant, 1996, p.112). Integral to knowledge work is technology, which shapes the practices through which workers construct, refine, and disseminate knowledge (Anthony, 2018; Bailey et al., 2010; Knorr Cetina, 1999). Commonly cited examples include drug discovery (Dougherty and Dunne, 2012); strategy making (Kaplan, 2011) and product development (Kellogg et al., 2006).

As developed societies transition from analogue to digital economies, firms are increasingly examining how to infuse algorithmic technologies into knowledge work to enable more efficient forms (Armour and Sako, 2020; McAfee and Brynjolfsson, 2012). Algorithmic technologies are computer-programmed procedures that can emulate human work and decision-making (Kellogg et al., 2020), and as such, are fundamentally different from traditional technologies (Bailey and Barley, 2020; Bailey et al., 2019; Benbya et al., 2021): they learn to make predictions from existing patterns in data (Agrawal et al., 2018), minimising the need for capturing tacit organisational knowledge into explicit formal rules. Prominent examples include machine learning (ML) and robotic process automation (RPA) algorithms. The ability to learn and, potentially, act, autonomously grants algorithmic technologies a different status from traditional technologies, extending the spectrum of roles that technology plays in knowledge work in general, and in improving knowledge work in particular.

The roles of algorithmic technologies in improving knowledge work can be understood from three perspectives. The first perspective is informed by an economic orientation, whereby algorithmic technologies are seen as a substitute for knowledge workers, affording the automation of a wider range of business activities (Arntz et al., 2016; Dellot et al., 2019; Manyika et al., 2017; Muro et al., 2019). Algorithmic technologies are posited as a 'deterministic' and independent force that acts upon knowledge work. That is, they are seen as a cornucopia of automation opportunities that lead to homogeneous effects in organisational properties such as structure and work. Frey and Osborne (2017), for example, argue that algorithmic technologies are a disruptive

force that exposes 47% of current employment to a high risk of being computerised. From his vantage point, improving knowledge work is enabled by reducing labour input or facilitating labour arbitrage – for instance by automating back-office (Lacity et al., 2016; Lawrence et al., 2016) and professional tasks (Sampson, 2021; Susskind and Susskind, 2016).

The second perspective is informed by a sociological orientation that highlights the role of algorithmic technologies as a reshaper of organisational structure and function in knowledge work contexts (Faraj and Pachidi, 2021; Faraj et al., 2018; Huysman, 2020). Building upon the classic contingency view (Thompson, 1967; Woodward, 1965) and the concept of socio-materiality (Barley, 1986; Orlikowski, 1992), scholars following this tradition argue that algorithmic technologies are endogenously co-constitutive of organising (Faraj and Pachidi, 2021). In other words, algorithmic technologies are entwined with organisational norms, rules, actors and infrastructure; and such entwinement enables or limits the emergence of new forms of organising knowledge work. This view stipulates that two organisations encountering the same algorithmic technology may realign roles and structures in different ways (Barley, 1990). Under these circumstances, organisations improve work when they use technology to align structure and function (Hayes and Wheelwright, 1979a, 1979b). In the case of knowledge work, for example, Sampson (2021) noted that algorithmic technologies may improve professional services by restructuring work, and augmenting some tasks, while deskilling, automating and centralising others.

The third perspective is informed by a learning orientation and underscores the role of algorithms as an inhibitor of learning (Balasubramanian et al., 2022; Choudhury et al., 2020). This perspective highlights that using algorithmic technologies may compromise firms' performance by reinforcing practices that prioritize short-term gains over long-term learning – what Levinthal and March (1993) referred to as 'myopic learning'. It emphasizes the propensity of knowledge workers to treat algorithmic recommendations as objective, disregarding the underlying uncertainty in knowledge inputs (Lebovitz et al., 2021) and the epistemic assumptions embedded in algorithmic predictions (Anthony, 2018). Furthermore, even when knowledge workers question

algorithmic recommendations, these recommendations can create impassable knowledge boundaries that heighten decision uncertainty, ultimately undermining the value of introducing algorithmic technologies in knowledge work. For instance, Lebovitz et al. (2022) showed that physicians experienced heightened uncertainty in diagnosis when algorithmic recommendations diverged from their professional judgement without providing an accompanying explanation. While organisations may deploy knowledge brokers to bridge these knowledge boundaries, unintended consequences can arise, such as brokers substituting algorithms' explanations with their own judgment (Waardenburg et al., 2021). These unintended consequences further diminish the usefulness of introducing algorithmic technologies in knowledge work.

4.2.2 A knowing-in-practice view of algorithmic technologies

Although research on the interplay between learning and algorithmic technologies constitutes an important step towards understanding the digital transformation of knowledge work, its current state falls short of adequately addressing the potential of algorithmic technologies to create operational improvement in knowledge work. To bridge this theoretical divide, we draw on the idea of knowing-in-practice, the point of which is to treat organisational knowledge (or rather, knowing) as an ongoing, situated, dynamic and embodied process that unfolds as individuals engage in organisational life through practice (Nicolini, 2011; Orlikowski, 2002; Tsoukas, 2005). Rather than being a static capability or an individual disposition, knowledge is viewed as a provisional achievement that is constantly shaped and reshaped through everyday practices (Hadjimichael and Tsoukas, 2019; Orlikowski, 2002; Østerlund and Carlile, 2005). By applying the concept of knowing-in-practice to the investigation of algorithmic technologies in knowledge work, we aim to explore how organisations effectively use algorithms to improve knowledge work.

To guide our exploration of how and when organisations can leverage algorithmic technologies to improve knowledge work, we focus on two aspects of knowing-in-practice. First, knowing in practice involves recognising that intelligently performing work relies on both the organisation's codified procedures (know-that) and the individual's

accumulated expertise (know-how) that is rooted in social and tacit practices (Lebovitz et al., 2021; Ryle, 1949). For instance, project managers know *that* conducting risk identification procedures is key for delivering projects on time and budget. However, executing this task intelligently requires project managers to possess the expertise to know *how* to assess and communicate potential risks, comprehend contextual uncertainties, and navigate power dynamics with project stakeholders. This process goes beyond rigid rule-following and involves integrating subsidiary (e.g., stakeholder's characteristics) and focal awareness (e.g., project risks) (Polanyi, 1967; Tsoukas, 2012). The interplay between knowing *that* and knowing *how* plays an important role in organisations' ability to harness the potential benefits derived from collaborating with algorithmic technologies in knowledge work. This points us to the challenges that organisations may face codifying situated knowledge into algorithms, often resulting in the development of algorithmic solutions that overlook critical contextual aspects and ultimately fail to yield desired performance improvements (Choudhury et al., 2020).

Second, adopting a knowing-in-practice perspective involves recognising the embodiment of knowledge within skilled performances (Hadjimichael and Tsoukas, 2019), and considering how this embodiment shapes the use of algorithms in knowledge work. Previous research has highlighted that knowing how to perform a task is ingrained in an individual's bodily skills (Dreyfus, 2014; Sandberg and Tsoukas, 2011). These skills are refined through ongoing practice, which involves "training our bodies to relate in certain ways to the world" (Tsoukas, 1996, p.17). Through repetitive practice, individuals internalize a set of background distinctions that constitute the work at hand (Hutchins, 1995; Lave and Wenger, 1991). This embodiment process includes assimilating a shared purpose, vocabulary, and set of values and norms that guide and legitimise their actions (Kellogg et al., 2006; Wenger, 1999). Recognizing the embodiment of knowledge in skilled performances has significant implications for technology adoption in general and algorithms in particular. For example, it suggests that skilled performances represent collective shared taken-for-granted ways of doing things embodied in individuals, which are difficult to capture in algorithmic technologies.

Scholars suggest that certain skilled performances, which rely on creativity, social intelligence, and physical manipulation pose challenges when it comes to 'computerising' knowledge into algorithms (Frey and Osborne, 2017). This idea is supported by the service operations management literature, where the need for creativity and social intelligence have been perceived as an obstacle to improvement in knowledge work (Lawrence et al., 2016; Lewis and Brown, 2012), not to mention its automation (Sampson, 2021). Creativity, the ability to generate novel and valuable ideas (Amabile, 1983), plays an important role in knowledge work. It allows individuals to understand and explore problems that involve variability in the inputs, processes or outputs (Schmenner, 1986). For instance, designing services tailored to client problems, or dealing with uncertain tasks requires creativity (Kellogg and Nie, 1995; March, 1991; Staats et al., 2011), especially in contexts involving making potentially risk-bearing decisions (Harvey, 2016; Schmenner, 1986). Similarly, social intelligence, which involves understanding and navigating social situations, including recognizing social cues, comprehending emotions and intentions, and responding appropriately in interpersonal contexts, is crucial for knowledge work performance in settings requiring high customer engagement (Brandon-Jones et al., 2016; Lewis and Brown, 2012); social intelligence allows individuals to diagnose customers' problems and suggest suitable solutions (Abbott, 1988; Schmenner, 1986; Silvestro et al., 1992). The quality of the solution heavily depends on the service provider's ability to interact and gather the necessary information to identify and understand the client's specific problem. Thus, in practice, creativity and social intelligence remain skills that are highly valuable but challenging to codify into algorithms.

4.2.3 Synthesis

Algorithmic technologies have been identified as objects of strategic consideration in knowledge work (Bailey and Barley, 2020; Faraj and Pachidi, 2021; Hollebeek et al., 2021; Sampson, 2021). Emerging important questions remain, however, regarding the relationship between algorithms and the improvement of knowledge work in practice (Brynjolfsson et al., 2019; Syed et al., 2020). Specifically, the conundrum that the

application of algorithmic technologies poses is that while the substitution of human capital for algorithms may reduce labour costs, at the same time it may harm the mechanisms by which organisations learn and produce value (Balasubramanian et al., 2022; Lebovitz et al., 2021). Building on a knowing-in-practice view, it is apparent that this conundrum is particularly relevant when algorithmic technologies substitute for skills in context, such as knowledge work, where tasks commonly require creative and social intelligence, and potentially also physical manipulation. In response, the underlying study explores the question of how knowledge workers experience, interpret and use algorithmic technologies to improve knowledge work.

4.3 Method

4.3.1 Research context and design

To investigate how organisations use algorithms to improve knowledge work, we conducted a two-year comparative multiple-case study within ACME, a multinational energy firm. ACME represents an ideal research context due to its reliance on professionals and experts engaged in knowledge work. Here, engineers, accountants and petrophysicists work together in business areas to create, share and refine knowledge that enables them to produce, trade and deliver energy products, including hydrocarbons, lubricants and low-carbon energy. Moreover, the context is appropriate because ACME was undergoing a digital transformation, with algorithmic technologies emerging as a strategic priority. The leadership team allocated US$1bn per year to capture, incubate and scale projects that used these technologies. As such, this context provided ample opportunities to gather data and elicit cases on how various business areas used algorithms to transform their knowledge work, and the consequences of this transformation for operational performance.

The case sampling logic we settled on was theoretical and comparative (Miles and Huberman, 1994), drawing inspiration from Davenport's (2015) four segments of knowledge work, which are ideal-type images of work that highlight differences in process orientation, as well as the degree of expertise and level of coordination required. The selected business areas with their respective images were business services (transactional), maintenance operations (integrative), internal audit (collaborative), and drilling operations (expert). For these business areas, we sought to trace the transformation of knowledge work as it unfolded over two years in seven different algorithmic technologies projects. This approach offered a strong foundation for meaningful comparisons while controlling for contextual variation based on matched similarities (Bechky and O'Mahony, 2015). The embeddedness of all business areas in ACME allowed us to control for factors that affect the adaptation to technological changes, including strategic orientation and technological aspiration (c.f. Eggers and Kaplan, 2009; Tripsas and Gavetti, 2000). ACME expected its business areas to participate in its digital transformation programme and adopt algorithms as productivity innovations. Furthermore, this embeddedness also ensured similarities between recruiting, governance principles and technology sourcing policies, allowing us to control for other known influential factors that may shape the adoption of algorithmic technologies, and hence exploit case variance to address our research question.

Business areas are a well-suited context for this study for several reasons. First, business areas are organisations where knowledge is particularly critical for producing services. Because fashioning services requires the interpretation of information and the construction of judgements, members of business areas must master technical and professional understandings that enable them to engage in practices for creating, analysing, transforming and warranting knowledge. Second, business areas provide a good

opportunity to generalise findings to a broader population, as they in effect operate like most medium-sized private service enterprises. Work is organised into tasks and conducted by groupings of specialists who must ensure their performance is financially viable and to the satisfaction of a wide range of customers. Third, business areas offer a context that involves salient tensions and dynamics in digital transformation. These tensions include, for example, balancing organisational, economic and institutional pressures to adopt algorithmic technologies with members' resistance to technological change. Thus, we argue that business areas offer a powerful site to study the digital transformation of knowledge work. Here, we briefly present the context of the four sampled business areas, laying the groundwork for this study.

Business services: Business services is a centralised area that provides business support services to the organisation, including credit management, customer service and human resources services. Because of the nature of services offered, members of business services have diverse backgrounds, typically, in finance, engineering, management and economics. Some members are also hired based on equivalent experience rather than education. Members of business services usually conduct standard work in the front and back offices. They work in large business process outsourcing centres and operate according to documented instructions compiled in standard operating procedures. For instance, in credit management services, members perform activities to assess the borrowers' ability to repay a loan according to the instructions in the accounts receivable standard operating procedure. While members of this area must be able to communicate effectively, their work does not depend on collaborating with others. Because work, in this case, was highly routinised and primarily governed by standard operating procedures, we argue that business services well matched the characteristics of Davenport's transactional knowledge work segment.

Maintenance operations: The Maintenance operations area provides asset management services to the organisation, including assessing the condition of machinery, developing maintenance procedures, and troubleshooting equipment issues. The members of this area, known as 'maintenance workers', are typically junior and mid-level specialists, technicians, schedulers, planners and coordinators. They are responsible for both planned and emergent work. Planned work involves regular maintenance activities to prevent unexpected asset breakdowns, while emergent work focuses on addressing initial faults to minimize downtime. Although maintenance workers have some discretion, they must also adhere to vendor recommendations and maintenance standards (e.g., ISO 55000 and 13374) to ensure compliance with legal, technical, and safety requirements.

Social intelligence skills play a crucial role in the maintenance business area. When a fault occurs, maintenance workers engage in communication with plant operators to gather information and diagnose the cause and location of the issue. This includes identifying potential exposure to corrosive, toxic, flammable, or other hazardous substances, as well as determining if the repair work needs to be conducted at heights, underwater, or in confined spaces. Collaboration and effective communication remain crucial during the actual repair work. Once a team of three to four maintenance workers is assigned to fix the fault, they communicate with the control room and other stakeholders to prevent the unintended restart of machinery around the affected area while maintenance workers are still working on it. A good understanding of social cues and trust among the team members themselves are vital to coordinate repair work and understanding emergent risk. Given the importance of standardized practices and collaboration in maintenance work, this case aligns with Davenport's integrational segment of work.

Internal audit: Internal audit is an area that provides independent financial and technical assurance services to the organisation, including verifying the appropriate operation of accounting processes, governance mechanisms and internal controls. Because internal audit offers both financial and technical assurance, members of this area (hereafter referred to as 'internal auditors') had different backgrounds, such as accounting and chemical engineering. The common denominator is that to become internal auditors, aspirants must master high levels of expertise; they must be skilled and seasoned professionals with a demonstrated command of regulations (e.g., the Sarbanes-Oxley Act), standards (e.g., the Generally Accepted Auditing Standards) and codes (e.g., the code of ethics). Such financial or technical prowess allows internal auditors to examine, verify and judge if the organisation is operating in line with mandatory legal guidelines and technical standards.

In addition to being subject matter experts, internal auditors rely on the mastery of collaboration skills to accomplish their job. Each year, internal auditors perform circa 200 audits spread over eight auditing rounds. In each round, a team of internal auditors (three to nine people) work together to write an audit's terms of reference and internal controls testing plans. They then collaborate with members of the auditee function (e.g., human resources) to execute such plans and produce reports of findings that outline risk areas and recommendations to improve the auditee's operations. Communicating effectively allows internal auditors to build trust with auditees and ensure the right evidence is collected and analysed, thus helping them ensure that the audit findings capture the auditee's operations in the best way. Whenever irregularities or risks to the organisation's growth or reputation are identified, internal auditors liaise with the chief auditor to convey them to the board of directors. For these characteristics, we considered this case as an exemplar of Davenport's collaborative segment of knowledge work.

Drilling operations: Drilling operations is an area that provides global drilling support services to the organisation, including defining the scope and activity of wells, overseeing the operation of wells, and designing fracturing and stimulation treatments. Members of this area (hereafter referred to as 'drilling engineers') are predominantly mid- and senior-level mechanical engineers, chemical engineers and petrophysicists who have an in-depth knowledge of drilling and completion operations. Before joining this area, candidates must demonstrate seasoned experience in dealing with a variety of practices and assets for drilling onshore and offshore wells; they must have at least five years' drilling experience, prove command of practices for drilling in subsea reservoirs, and understand managed pressure drilling equipment – specialised arrangements of rotating control devices, choke manifolds and other apparatuses that control the pressure in the wellbore. Their extensive experience grants drilling engineers a high level of discretion. Some expert senior drilling engineers, for example, are responsible for designing tailored fracturing and stimulation treatments for oil and gas wells with unconventional environmental conditions around the globe. These designs depend on acumen, as drilling engineers draw on past experience with similar wells to analyse technical and regulatory dimensions that may impact the safety and cost-effective operations of the present well. For these characteristics, we matched this case with Davenport's expert segment of work.

4.3.2 Data collection

Data collection at ACME began in November 2019 and finished two years later. We acquired access to the four business areas outlined above and collected data using multiple methods, including interviews, participative observations and archival data, as summarised in Table 4-1.

We conducted 55 semi-structured interviews with members of the business service (n=18), maintenance operations (n=11), internal audit (n=15), and drilling operations

(n=11) areas. These interviews averaged 75 minutes each and were carried out individually in private rooms or via Zoom. Our informants had diverse backgrounds, covering engineering, accounting and data science, and included members with and without leadership responsibilities. Informants were invited to participate in the study because they drew mainly on knowledge and information to create value for each area, and because they participated in projects that used algorithmic technologies to transform their work. This allowed us to get '*retrospective and real-time accounts by those people experiencing the phenomenon of theoretical interest*' (Gioia et al., 2013, p.19). To ensure the validity of our findings, interviews were conducted until we achieved theoretical saturation (Glaser and Strauss, 2017).

Initially, we asked informants about their everyday tasks and how knowledge was created, used or shared. Our interview protocol included questions focused both on quantifying how much creativity, social intelligence and physical manipulation were required to complete the task (on a 10-point Likert scale); and on qualifying the role of knowledge in work by collecting examples that supported their rankings (see Appendix B). For example, an auditor argued that the task '*Test the design and operating effectiveness of controls*' involved a high degree (8 out of 10) of social intelligence, as it required negotiation and collaboration with auditees and experts from different occupations. To understand how informants perceived opportunities and/or resisted adopting algorithmic technologies, for each task the informant described, we asked a question on the perceived fit of algorithmic technologies to automate the given task.

As the research progressed, our interview questions focused on how each business area used algorithms to improve knowledge work processes. Participants were asked to describe automation projects that used algorithmic technologies, their scope (which tasks were affected) and consequences (benefits and drawbacks experienced after using

algorithms) (see Appendix C). We phrased our questions carefully to avoid leading interviewees: rather than asking them to contrast the effect of different projects involving algorithmic technologies, we asked them '*Tell me about a positive experience automating work with algorithms*', and '*Tell me about a negative experience automating work with algorithms*'.

To complement our interviews, we drew on participant observations. We visited each area multiple times between November 2019 and March 2020. These visits took the format of one-week participant observations, apart from one COVID-related exception. The first author was given a desk, training on the field, and permission to contact anyone in the organisation. We participated in day-to-day knowledge work activities, such as maintenance planning meetings, in which engineers discussed short- and long-term asset management projects to ensure complex plants run on time; or internal audits, where auditors interviewed corporate employees in, for example, human resources, to identify potential risks and recommend how to address them. These observations were designed to complement our interviews by capturing aspects that participants may have struggled to articulate during interviews (Barley and Kunda, 2001). Thus, we used them to create 'thick' descriptions of how participants draw on tacit knowledge to practice (Bechky, 2006; Orr, 1996). The focus of observations was twofold: to capture how the areas adapted knowledge work around algorithmic technologies, and to build a more comprehensive picture of how professionals and experts create and share knowledge before and after infusing algorithms at work. We made our research overt, and maintained a journal with notes about the business area's practice, keeping a separate journal for recording experiences, anxieties and reflections from fieldwork to minimise contaminating data with biases (De Rond and Lok, 2016). By visiting each area multiple times, we built a rapport with participants and minimised the risk of reactivity.

After the national lockdown was instigated in March 2020, we paused the observations and began collecting archival data as ACME operated a work-from-home policy. We gathered standard operating procedures that governed the practices in each of the four business areas and internal communications that captured the impact of algorithmic projects on the company's operational performance. These documents helped us understand the effects of algorithms in the four segments of knowledge work.

Table 4-1 Description of cases and data collection

Business Areas Characteristics						Data Collection		
Area	**Knowledge work segment**	**Degree of expertise**	**Degree of collaboration**	**Work characteristics**	**Examples of tasks in this case**	**Informants**	**Interviews (n=55)**	**Additional data collected**
Business services	Transactional work	Low	Low	Work is repeatable, reliant on formal rules and easily structured in process terms. Workers have low discretion	Issue credit cards for customers Process clients' procurement orders Input purchase orders into ERP systems	Data scientists (9) Automation group director (1)	18	
Maintenance operations	Integrative work	Low	High	Work is repeatable, reliant on formal rules and easily structure in process terms. Workers rely on tight integration across functional boundaries	Prepare asset maintenance plans Collect data about assets' performance Prepare detailed reports on maintenance effectiveness	Maintenance workers (9) Functional Manager (1) Deputy Manager (1)	11	Observational notes
Internal audit	Collaborative work	High	High	Work is improvisational and judgement oriented. Workers rely on tight integration across functional boundaries	Test the design and operating effectiveness of controls Prepare detailed reports on audit findings Plan the audit terms of reference for compliance-based and risk-based audits	Financial Auditors (4) Operational Auditors (5) Data Scientists (2) Audit Managers (1) Functional Manager (2)	15	Archival documents (standard procedures, press releases, industry reports)
Drilling operations	Expert work	High	Low	Work is improvisational, judgement oriented, and dependent on star performers	Design fracturing and stimulation treatments for offshore wells Design fracturing and stimulation treatments for onshore wells Prepare and manage the completion team to ensure the acidising treatments are done effectively	Drilling Engineers (6) Data Scientists (2) Functional Manager (1)	11	

4.3.3 Data analysis

Our analysis was inductive, iterative and comparative. The aim was to elaborate theory on the adoption of algorithmic technologies and transformation of knowledge work by discerning patterns in and across cases. As the foundation for this analytic strategy, we used the constant comparative method (Glaser and Strauss, 2017). Over various iterations, we read and coded the interview transcripts and observational notes, looking out for similarities and differences in the adoption of algorithmic technologies across the four business areas. These iterations took place over three main stages of analysis: understanding current knowledge work, identifying the knowledge work transformation process triggered by algorithmic technologies, and explaining differences in outcomes in the adoption of algorithmic technologies. These stages are detailed below.

Stage 1: Understanding current knowledge work. In the first stage, we analysed the interview transcripts and field notes to build narratives that described how the four business areas enacted knowledge work before encountering algorithmic technologies, helping us understand '*how things work*' (Wolcott, 1994, pp.12, 16). Using these narratives, we listed every task performed by members of each business area, considering a task a '*unit of work activity that produces output (goods and services)*' (Acemoglu and Autor, 2011, p.1045). Examples of tasks coded included 'processing clients' purchasing orders', 'planning the terms of reference of an audit' and 'investigating issues on fracturing and stimulation'. By grounding our task list on the narratives and practices observed across the four areas, we captured *variation* in work otherwise obscured in work databases (Bailey and Barley, 2020). For example, we coded narratives on 'planning an audit's terms of reference' into two distinct tasks, recognizing that auditors performed different activities and required different skills based on whether the planning involved compliance-based or risk-based audits. This task list provided a

foundation for analysing subsequent changes brought about by adopting algorithmic technologies.

We also coded each task in our list using five dimensions that measured the degree of creative intelligence, social intelligence, physical manipulation, repetition and perceived suitability for automation. These categories were grouped to form the intermediate basis for data analysis. The purpose of this phase was not to provide a quantification of the effect of the nature of the task on knowledge work automation, but rather to observe whether any meaningful pattern emerged between each dimension and the applicability of algorithmic technologies. We used this to deepen our analysis of the case studies, seeking to identify work characteristics that enable the use of algorithmic technologies in knowledge work. At this point, we sent a case write-up to the business areas' senior directors to validate and refine the cases. Directors had more than 10 years' experience in the job and understood in detail its areas' processes. They provided feedback and suggested some other people to contact to bridge the gap in data coding and interpretation (Hesse-Biber and Leavy, 2011).

Stage 2: Identifying the knowledge work transformation process triggered by algorithmic technologies. In the second stage of the data analysis, we focused on understanding the process of transforming knowledge work. To that end, we examined a subset of successful and unsuccessful projects that involved algorithmic technologies. We identified seven projects across the four business areas that included the application of machine learning, robotic process automation and logic-based algorithms, which are summarised in Table 4-2. For each of these projects, we then interviewed the professionals and experts involved, the developers who built the algorithm project, and the managers who sponsored and oversaw it. To build up descriptive abstractions from interview notes and field observations on algorithm

projects, we used NVivo and applied the constant comparative method, following the guidelines suggested by Gioia et al. (2013). Because we assumed that informants were '*knowledgeable agents*' who could describe the implementation of algorithmic projects as well as explain resulting changes in their work, we allowed first-order codes to emerge from quotations and observational notes. This initial coding followed a line-by-line approach (Charmaz, 2006). Then, applying axial coding (Corbin and Strauss, 2008), we built up descriptive theoretical categories that captured similarities and relations in first-order codes. The final step of this analysis involved distilling and mapping the relationship between these themes to develop four dimensions that explained the transformation of knowledge work. The result of this second analysis stage was a grounded theory model of the adoption of algorithmic technologies in knowledge work. The resulting data structure that served as the basis for our findings is presented in Appendix C.

Table 4-2 Descriptions of projects involving algorithmic technologies across cases

Area and Project	Project Description	Status
1. Business services: Credit card onboarding	An effort to build a robotic process automation solution to automate the credit card onboarding process.	In operation
2. Business services: Credit card locking	An effort to build an intelligent process automation solution to automate the decisions of locking or unlocking customers' credit cards.	In operation
3. Maintenance operations: Maintenance scheduler	An effort to use logic-based algorithms and ranking systems for generating and updating preventive maintenance schedules.	In operation
4. Internal audit: Anomalies in financial transactions	An effort to build a solution to automate the testing of operating effectiveness controls in compliance-based audits by using rules-based and machine learning algorithms for discovering anomalies in financial transactions.	In operation
5. Internal audit: Risk reports for directors	An effort to create a machine learning solution (topic modelling and Latent Dirichlet Allocation) that recognised natural language from audit reports to prepare quarterly summaries for ACME's directors.	Cancelled
6. Drilling operations: Drilling optimisation	An effort to develop a reinforcement learning agent to automate drilling onshore and offshore.	Paused
7. Drilling operations: Mini frac calibration	An effort to build a reinforcement learning agent for calibrating a 'mini frac' on an offshore well. (A 'mini frac' is a small fracturing procedure to observe the behaviour of the fracture system, and to calibrate the main hydraulic and stimulation treatment.)	In operation

Stage 3: Explaining differences in outcomes in the adoption of algorithmic technologies. In the final data analysis stage, we explored the variance across cases associated with each dimension of the transformation of knowledge work process identified in the previous analysis stage. Our focus was on understanding dynamics that may affect the interpretation and use of algorithmic technologies. We began by tracing how the decision to deploy algorithms unfolded in each project, and subsequently investigated the pathways of transformation that occurred between the decision to launch an algorithm project, and the emergence of the new work practice. Because of our interest in knowing in practice, we focused on the characteristics and aspects of knowledge work that enabled its transformation. Finally, we compared how each project resulted in adjusting the knowledge regime. Although it would also have been interesting to examine the role of organisational actors in the transformation of knowledge work, such an analysis lay beyond the scope of this study.

4.4 Findings

Our analysis of how the transformation of knowledge work unfolded across the seven projects led us to two important findings. First, we found a loosely five-stage transformation process that was common to all cases. This process began with a stage of 'overcoming inertia' through which all business areas navigated a constellation of interests, capabilities, powers and needs to decide whether to deploy algorithms and to what extent. Once a project was agreed upon, the transformation of work proceeded in either of two pathways: one in which the business area focused on 'altering a specific task', and where actions were taken to transform a unit of work undertaken typically by individuals or small groups; and another in which the focus was on 'reengineering the entire process', where actions were taken to transform the bundles of activities that

convert inputs to outputs. Our analysis uncovered that these pathways lead to a stage of 'adjusting the knowledge regime' whereby some business areas redefined their structure and authority arrangements, as well as their schemes for evaluating knowledge.

Second, while in each of the seven projects organisations underwent a similar transformation process, a closer examination revealed differences in how the business areas experienced, interpreted and used algorithms to transform their practice. Evidence supporting such variation is presented in Tables 4-3, 4-4 and 4-5. We use these variations to theorise, for example, under what circumstances organisations follow the altering of the task pathway, and when they reengineer the process pathway; what affects the ability to use algorithms to improve knowledge work; and how organisations can break the path dependency induced by technological transformation. In the remainder of this section, we discuss our findings concerning five broader themes that emerged from our fieldwork, before summarising these findings in a conceptual model.

4.4.1 Overcoming inertia

Our first set of findings concerned how organisations broke organisational inertia to transform their practice using algorithmic technologies. The decision to adopt algorithms in the four business areas, in common with many organisations, did not always follow a rational decision-making process that sought to optimise the match between technology and work. Instead, we found that such a process followed an imperfect, affective and political rationale that varied in the kinds of attitudes towards and interpretation of algorithmic technologies. When most members of an area had a positive attitude towards algorithms, and perceived algorithms as opportunities, the adoption of algorithms followed a logic of *coping* or *differentiating*. When most members of an area held a negative attitude towards algorithms and perceived them as threats, adopting algorithms followed a logic of *coping* or *conciliating*. Finally, adopting algorithms followed a logic of

challenging when most members of an area held a negative attitude towards algorithms, but interpreted them as opportunities and threats simultaneously. In the remainder of this section, we discuss how the attitudes and interpretations of algorithms led to each of the four adoption logics.

4.4.2 Transforming knowledge work

Our second and third set of findings concerns how organisations used algorithmic technologies to transform knowledge work. We found that organisations followed two pathways: one for altering specific tasks, and one for re-engineering an entire process. Because of our interest in knowing in practice, we focused on understanding how organisations dealt with the nature of work to create opportunities for using algorithms in knowledge work.

Our findings show that social intelligence and creativity were obstacles to using algorithmic technologies in knowledge work, which aligns precisely with recent studies (Balasubramanian et al., 2022; Frey and Osborne, 2017). To understand this observation in more detail, we characterised the suitability of algorithmic technologies for different tasks, by asking knowledge workers to rank in which tasks algorithms were applicable. We then matched these ranks against practitioners' consensus on the degree of creativity, social intelligence and physical manipulation required to perform a task. These were valid measures for our analysis, because, for example, for a task to be creative, other experts must judge it to be creative. These judgements are not absolute (i.e., a solution is creative, or not creative), and reflect a degree of skill required to perform a task (see the methods section for more details on the analysis). The result of this characterisation is presented in Figure 4-1. This figure reveals a more differentiated view of Frey and

Osborne's (2017) findings regarding how the skills required to perform a knowledge work task affect the ability to computerise it.

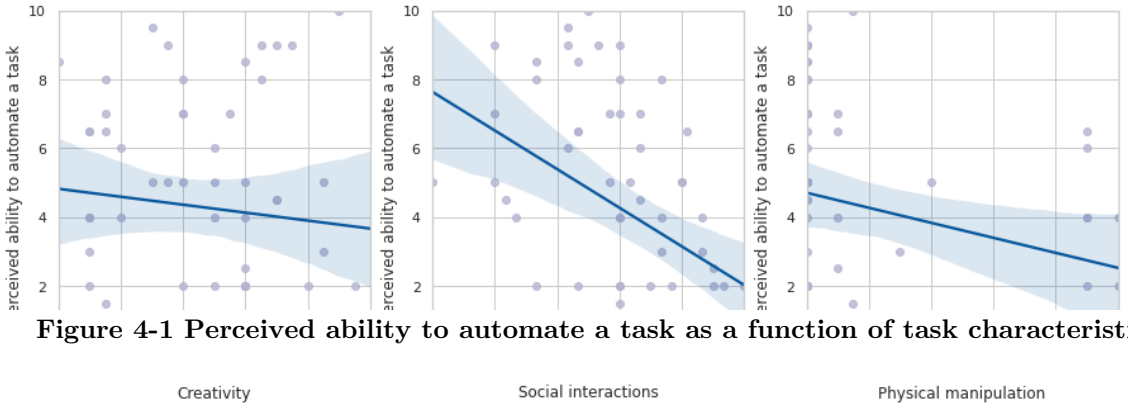


Figure 4-1 Perceived ability to automate a task as a function of task characteristics

In a broader sense, our findings suggest that automating creativity seems to be less of a hurdle than automating social intelligence, which is consistent with what emergent studies on the automation of professional services observed (c.f. Sampson, 2021). By creativity, we mean work that involves an experimental activity in which actors (humans or not) discover a problem and then develop a previously unclear path to a solution, which addresses the said problem (Amabile, 1983). By social intelligence tasks, we mean tasks that involve being empathic to people's needs and feelings, that require persuasion or negotiation, or that depend upon being personally connected to others – as specified in the O*Net database. Regarding automating physical manipulation – tasks that involved the coordination of movements of the fingers or hands to manipulate and assemble objects (Frey and Osborne, 2017) – the results were less clear, since most knowledge work tasks studied required no manual dexterity. In the few cases when we observed work that required physical manipulation, we found that organisations indeed struggled to meaningfully use algorithms in such work.

Yet, our findings point to a more important difference from previous research (c.f. Arntz et al., 2016; Frey and Osborne, 2017). Namely, we observed that at the organisational level, there is a high degree of heterogeneity among the task characteristics and the use of algorithms to automate it. This level of heterogeneity was also identified

during the later analysis of the automation projects – which we will discuss shortly – suggesting that these skills alone are insufficient to explain the ability to use algorithmic technologies to transform knowledge work. In particular, Figure 4.1 suggests some tasks require a high degree of creativity and social intelligence, where algorithms are still applicable. This finding suggests that organisations use practices to work around the need for 'uncodifiable skills'.

Using participants' accounts and an in-depth analysis of seven algorithm projects, we now explore these practices that allow organisations to use algorithms in work that requires social intelligence, creative intelligence or physical manipulation. We will discuss these findings in relation to the transformation pathway that organisations followed: altering specific tasks, or re-engineering entire processes.

Altering the task pathway

The first pathway of knowledge work transformation focused on altering specific tasks. This pathway was found in four projects and occurred mainly in internal audit and maintenance operation areas – those with a defensive stand. In each case, efforts to transform the practice were directed towards units of work done, on questions of skills, repeatability and standardisation. Members of these areas, for example, often sought to triage tasks into common and uncommon cases, enforce standards in activities that require creative intelligence, and divide work into cognitive and physical labour. By performing these practices, organisations aimed to use algorithms to achieve task efficiency gains. Those transforming work by altering the task engaged in practices such as *triaging* and *standardising* tasks to enable the use of algorithms in this work. We discuss these practices below in Table 4-4.

Table 4-3 Transforming work by altering specific tasks

Area and project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
1. Maintenance operations: Maintenance scheduler	Every week, the maintenance scheduler creates a document outlining the maintenance activities needed for the upcoming week. This document defines the required maintenance work (e.g., replacing needle valves), its priority, the deadline, and assigns maintenance workers. To this end, the scheduler assembles preventive and corrective maintenance activities into a spreadsheet. The creation of the spreadsheet involves pouring over asset manuals and maintenance logs to identify what parts need to be replaced to prevent future faults. Also, it involves interacting and collaborating with various internal and external members. The scheduler communicates with maintenance workers to identify emergent faults, with managers to check what maintenance workers are capable and available for a given task, and with members of the operations team to ensure shutting down an asset on a given date would not impact production. Because emergent issues arise constantly and some spare parts may be unavailable, the maintenance worker revises the spreadsheet several times. The turnaround for this task is about four days.	As the maintenance scheduler did not have the management support to invest in new technologies, she repurposed her tools by using writing macros in Excel VBA. As she wrote the code, she standardised some metrics, for example, to rank maintenance tasks. She showed the tool in confidence to the maintenance planner. They agreed to test and debug the tool before revealing it to the rest of the team. Debugging happened over multiple iterations, and was continuous and informal.	This work resulted in transforming the way maintenance schedules were produced. The approach was atomistic and focused on ranking maintenance activities. The maintenance scheduler remains in control of activities that require social intelligence. She gathers information and inputs it into the system, where the algorithm recalculates the maintenance activities that need to be done that week.
2. Internal audit: Inconsistencies in financial transactions	During every internal financial audit, internal auditors collaborate with members of the auditee area to gather and analyse accounting records, checking for inconsistencies, regulatory compliance issues, and lack of adherence to accounting best practices. This	Auditors with experience in data science were hired as brokers of information. They were assigned to lead a new team	This project resulted in altering the evidence collection task. The new practice requires auditors to collaborate with data scientists to gather and test evidence of financial transactions. Here, an auditing team composed of

Area and project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
	work includes applying sampling and testing techniques, tracing the amounts and disclosures in financial statements, running enquiries and conducting third-party confirmations. Social intelligence is crucial to ensure a smooth audit. Internal auditors collaborate and build trust with auditees, for example, to identify relevant evidence and investigate incidents in detail. Social intelligence also allows internal auditors to pick verbal and non-verbal cues, such as if an auditee has a lax attitude or feels uncomfortable discussing a topic – indications that increase the risks of financial mistakes and suggest the need for supplementary investigation.	within internal auditing focused exclusively on creating scripts for gathering financial records and testing them for mistakes.	auditors and data scientists is ensembled when an audit's terms of reference include the inspection of financial transactions. The audit lead specifies to the data scientist what evidence needs to be collected and tested. The evidence is then inspected by the rest of the auditing team, who set recommendations for a given auditee function, allowing auditors to remain in control of their occupation jurisdiction.
3. Internal audit: Risk reports for directors	At the end of every quarter, internal auditors present to ACME directors a summary of key audit findings and identified risks. To prepare the summary, a team of internal auditors work together to analyse and synthesise 50 audit reports. This task is social in nature; during the preparation of the summary, internal auditors discuss and make sense of how diverse operational and non-operational risks may aggregate across business units. Furthermore, this task requires knowledge translation; to create risk summaries accessible and interpretable to directors, internal auditors must first separate internal audit report findings from their original context – isolating jargon and descriptions of evidence, gathering and testing – and translate them into PowerPoint slides with figures and accompanying narratives that	Internal audit delegated the project to ACF. ACF asked internal auditors for all the previous year's audit reports, as well as the equivalent report findings summaries presented to auditors. ACF used this data to identify what the algorithm needed to learn. The training model seems to match state-of-the-art algorithms for NLP. Internal audit assumed that because ACF has major experience in auditing, they	The team trained a natural-language-processing algorithm to automate this task, yet results were disappointing. They partnered with one of the big four accounting firms, ACF, whose team of expert data scientists trained the algorithm to discover abstract topics occurring across reports; and provided them with a historical collection of ACME's audit reports and granular descriptions of audit findings. Yet, after a few months' work, the team discovered that the results did not make sense from a business perspective. One senior auditor found that the algorithm flagged several risk areas that were completely outside the scope of the audit reports. 'There was no correlation between what the audits covered and what the tool showed', he noted. The

Area and project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
	aim to focus directors' attention to specific risks.	will know how the internal processes are run within ACME.	algorithm produced summary reports that flagged, say, cybersecurity as a key threat to ACME, while no cybersecurity issue appeared in the original audit reports. The Internal Audit team was concerned that these inaccurate summaries could lead to the redirection of management efforts into ineffective actions. The audit organisation thus decided to cancel the project.

Triaging tasks

One practice adopted by those transforming knowledge work by altering specific tasks involved triaging tasks. By triaging tasks, we mean the practice of separating tasks, either formally or not, based on whether they involve common or uncommon situations. For example, organisations discriminated between tasks requiring social intelligence for information gathering from known actors, formats and systems, and social intelligence for untangling '*…ill-defined and emergent situations*'. While both cases relied on trust and effective communication to gather and manipulate information, participants agreed that the former applications were familiar and virtually constant, so no real intelligence was required. The classic example was gathering information from clients and colleagues to populate fields in customer relationship management software (e.g., Salesforce).

Triaging tasks into emergent and common cases became a powerful way to identify where algorithms could lead to opportunities to improve efficiency, and instances for computerising tasks performed in what participants described as a '*quite old school way… still relying on a human interface*'. One participant mentioned that this strategy allowed them to use algorithms to '*integrate the flow of information through different systems,*

collating and manipulating and validating [knowledge], which helped [them] to really work quickly'.

Standardising tasks

Organisations transforming knowledge work by altering specific tasks also worked around requirements for skilled performances by standardising tasks. This practice allowed them to constrain the amount of creativity to complete the task by capturing the know-how of the actual practice into standard operating procedures. This subgroup was supportive of ACME's organisational learning activities, engaging with both internal and external members in workshops to capture best practices in the industry. Consider, for example, the case of knowledge workers in the internal audit area. Internal auditors adopted a continuous improvement mindset to capture lessons learnt and standardise their testing methodology. They further partnered with ACF, their external auditing firm, to redesign ACME's audit practice concerning what other companies were doing. They argued that standardising their approach to auditing allowed them to '*… have consistency for testing'* (Senior Auditor); '*…ensure that we test everybody in the same way so that we've got benchmarks that we can review*' (Audit Manager). These efforts to capture external and internal knowledge into their practice ultimately pave the way for using algorithms in the auditing practice. Internal auditors substituted the need for creativity in designing audit test plans with standard test plans, enabling them to use, for example, algorithms for document automation. '*We're not doing everything differently every time*' (Senior Auditor); '*We might make some tweaks and modifications to it, depending on the nature of the area we're looking at or what the focus is… But at the core of it, it's a standardised testing approach.*'

Separating 'hand and brain'

A final practice that served organisations to alter specific tasks was separating 'hand and brain'. Although ACME struggled with the automation of tasks that required both physical and knowledge manipulation, we found that organisations separated work into hand (manual work) and brain (cognitive work) to enable the use of algorithmic technologies in the latter. Unlike the past separation of hand and brain, which focused only on the technification of physical work, we observed a division of labour that aimed at using algorithms in decision-making. We noted, for example, that the drilling business area underwent a separation of the work of fracturing experts who did mental work (those responsible for designing fracturing and stimulation treatments) and physical work (those responsible for wells completion). One engineer mentioned that the new division of labour that the organisation was undergoing, had a different goal: it aimed to minimise the need for knowledge labour. They set up sensors, algorithms and protocols so that fracturing specialists could work remotely instead of being in the field. One engineer shared: '*We use satellites to get this real-time data… we can look at 30, 40 different channels of data streaming into one pressure curve plot. You can customise the plot to what you want, automate decisions, and make real-time communication with an offshore facility or a remote facility.*' This further separation between hand and brain led to the centralisation of knowledge workers, and their isolation from physical labour.

Re-engineering the process pathway

The second pathway of transformation focused on re-engineering the process. This pathway was found in four of seven projects, occurring mainly in the business services and drilling engineering areas. In each of these cases, the focus was on 'processes' – bundles of activities that convert inputs to outputs. That is, organisations following this pathway focused on redesigning the entire process to enable the automation of some of its embedded tasks. Members of these areas were concerned less with the question of

skills and more about reconfiguring the way the organisation transformed resources into services. Some practices deployed by those following this pathway included controlling for process variation, adding steps for validating the co-construction of knowledge, and eliminating non-value-adding activities. Below, in Table 4-4, we describe these dynamics in more detail.

Table 4-4 Transforming knowledge work by re-engineering the entire processes

Area and Project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
4. Business services: Credit card onboarding	At a high level, this process starts when a customer puts in an application for a credit card. A team of knowledge workers first run various customer due diligence checks across different systems – credit systems, enterprise resource planning systems and customer relationship systems –to decide whether to approve or reject a credit card application. While some information-gathering activities had been automated, others still required members to schedule calls and gather information manually. Once all information is gathered, knowledge workers decide to approve or reject the customer application. If the application is approved, they issue a physical card to users, who can start using it after activating it. A senior developer shared with us: 'So this whole process was a very integrated process across multiple applications, and it was taking a lot of time. And many times, we were not able to issue the cards to the customer till at least around a month, which basically means that if people are interested to take [the card], obviously they have to wait.' The turnarounds to onboard a customer were about 30 days.	The group partnered with the CoE to design and deploy the algorithm solution. The CoE hired experts on process mapping, who interviewed various stakeholders in the credit card onboarding process to identify steps for issuing a new card. Formal validation of this process mapping followed. The CoE sent reports to product owners to ensure the process was appropriately captured. The team removed non-value-adding activities and artefacts. Key to this removal was challenging the assumption that customers needed a physical credit card to operate. They then adjusted internal policies to allow customers to use virtual credit cards. A dashboard was created to monitor the performance of the RPA algorithms over time.	A new credit card onboarding process was created. The new process triages credit card applications into common (processed by RPA and logic-based algorithms) and uncommon cases (processed manually). The new process issues a virtual credit card, challenging the assumption that customers needed a physical card to use the credit product. Turnaround is considerably decreased.

Area and Project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
5. Business services: Credit card locking	Business service analysts lock customers' credit cards for various reasons, such as suspected fraud, unusual spending patterns, high-credit utilisation and late payments. To do so, analysts manually inspected databases, checking if customer payments were on time and if there were any suspicious transactions. Most of these inspections were manual, and in some cases, analysts needed to schedule calls with database owners to gather financial information. This activity required negotiation and interpersonal skills to get database owners to prioritise analysts' requests over other work. This helps analysts lock and unlock credit cards as promptly as possible. Despite these efforts, the turnaround for locking and unlocking customers was significant, increasing the risk for business services.	As before, the business services area collaborated with the CoE to create the algorithm solution. The new solution uses robotic process automation. To understand and control the variation in the credit card locking process, the CoE interviewed process owners and created entity relationship and process maps. Using these boundary objects, the CoE created multiple RPA algorithms that dealt with variations in credit card locking systems. To monitor the performance of algorithms, the CoE created a dashboard that allows team members to validate that the algorithms are performing according to specifications. A person is assigned to monitor continuously the logs generated by the algorithm, including exceptions that may indicate something is wrong.	As with the previous case, the project involved the redesign of the credit card locking and unlocking process. This project, however, was more difficult to execute as there were a lot of variations; requests could come from different sources, and checks required the management of different systems. Adding more logic to cope with these requests increased the complexity of the tool and also its maintenance.
6. Drilling operations: Drilling optimisation	The drilling process starts after the site has been prepared with the relevant infrastructure and the permits to drill have been secured. The process involves drilling a hole, called the wellbore, to the desired depth; logging measures and samples to determine the productivity of the reservoir; and completing the wellbore (casing, and tubing to prevent the well from collapsing	To execute this project, the drilling operations area hired PhDs who had experience building machine learning algorithms and who understood, at least theoretically, the geological and hydraulic mechanisms involved in drilling. These PhDs acted as	A reinforcement learning algorithm runs drilling simulations to find the parameters that yield the 'optimal' drilling speed while ensuring temperature and pressure remain on a prespecified

Area and Project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
	and to allow oil and gas to flow to the surface). Drilling a wellbore can be done on land-based or offshore rigs. While these two modes of drilling involved different levels of expertise, both require expert workers. These workers need to understand the geological principles of rock formation and fluid content of the wellbore, be able to design and oversee drilling processes, and understand the mechanical and hydraulic and electrical systems involved in drilling equipment.	boundary spanners. They wrote python scripts to demonstrate the potential of algorithms in drilling. Their focus was on compressing the drilling process, by removing unnecessary work and using algorithms to automate others. Various metrics were used to measure the validity of algorithmic recommendations in drilling, such as the coefficient of determination. The team also created an interface to allow the operations control team to validate in real-time if algorithms were performing appropriately.	and safe window. Despite showing potential, the project is currently paused because the team found variations in the simulation time that render the algorithm ineffective. Drilling engineers deemed that the best solution to control such variations is to bring the simulation environment into their ecosystem. In this way, they can remove data handshakes.
7. Drilling operations: Alpha 2: 'mini frac' calibration	Work affected in this project involved the calibration of offshore hydraulic treatments – fracturing geological formations to stimulate oil and gas production. Unlike fracturing onshore, fracturing offshore is considered creatively complex by most experts. The team must design tailored fracturing solutions and adjust them to react to unconventional conditions in the well and the environment. For one, they conduct operations on subsea reservoirs that are 25,000 feet beneath the sea and, in their designs, they must consider the effects of storms that sometimes batter the vessels in which the operations are carried. If the process goes wrong, it cost ACME millions – a situation that discomfits even the most veteran specialists. Because drilling	Drilling engineers removed non-value-added activities, compressing the process of fracturing and stimulating wells, and simplifying the methodology to drill. They described this as a shift to drive their work 'towards simplicity and towards flexibility' (Senior Engineer). This simplicity, in turn, enabled the use of algorithmic technologies. During this shift, they questioned the need for using traditional waterfall methods for managing drilling projects and moved	The result of this project was a new process in which knowledge workers collaborate with reinforcement learning algorithms to design fracturing and stimulation treatments. The algorithm allowed the team to improve substantially the time to perform the mini frac (for reasons of confidentiality, we refrain from mentioning the size of the improvement). Eventually,

Area and Project	Nature of Work before Algorithm Adoption	Practices for Transforming Work	Outcome
	engineers cannot see directly what happens inside the well, specialists apply a set of diagnostic techniques to characterise the reservoir, including the so-called 'mini fracs' to observe the behaviour of the fracture system and decide how to calibrate the main fracturing treatment. They accompany this observation-diagnosis process by revisiting their past experiences, running scenarios to create mental contingency plans.	towards using agile methodologies.	drilling engineers questioned the need for doing a mini frac and decided to merge it with the main fracturing and stimulation treatment activities.

Identifying and controlling process variation

One of the practices involved in re-engineering the processes was related to identifying and controlling the process variation. Process variations occurred when the conditions and attributes of a process varied, preventing the process from following a constant pattern. These variations were important because they hampered the applicability of algorithms in work. A data scientist working in the business services area mentioned in an interview that processes that '*(…) have a lot of different rules to go through and variations make the logic more complicated, and that's where you could struggle with the success rate of technology*' (Senior Developer). These knowledge workers were willing to do some pre-work with users to understand variations in inputs (e.g., if the knowledge worker interacted with the clients via multiple channels), conversion processes (e.g., if the logic to make decisions varied widely across knowledge workers), and output (e.g., if the resulting product or service was tailored to a client's needs). By identifying sources of variation, they pinpointed the mechanisms to deploy to control such variation before using algorithmic technologies.

A good example of this practice involves a project in the business services area to incorporate algorithmic technologies in locking and unlocking credit cards based on

customers' transactions, payments and credit information across platforms. At first glance, the task seemed like simple automation of data collection. But soon, the business services area realised that the activity had many potential variations and interdependences that risked the project's success. Particularly, the algorithm logic became complicated as credit card payments came from different intermediary banks, which used different data formats and processing times. Variation further increased as developers added non-payment-related transactional data and rules to minimise defaulting risk. One senior data scientist noted, '*[Locking and unlocking credit cards] has been very tricky. It has a lot of dependencies on all sort of different [payment] processes… it was quite complicated in terms of looking at all aspects before you make a decision to lock a customer.*' To deal with the inherent variation in locking or unlocking customers' credit cards, data scientists interviewed and conducted workshops with owners of the credit card process to map tasks, inputs and dependencies involved in locking and unlocking credit cards. The results of these workshops included entity relationship models and process maps that acted as boundary objects to surface potential variables and variations in activities that might influence the decision of locking or unlocking a customer. These boundary objects, in turn, allowed business services to create more complex algorithm solutions that bridge identified process variations, reducing service times to lock and unlock credit cards in real time.

Adding steps for assuring output quality

Another practice involved in re-engineering the process revolved around assuring output quality. As the name indicates, this practice entailed redesigning work by adding extra steps to the process to assure the quality of knowledge created by algorithms. Consider, for example, the project to use reinforcement learning to automate the calibration of mini fracs in offshore wells – a task considered to require high levels of creative thinking that

involved characterising a well by causing a short fracture in the formation before the main fracturing treatment (see Table 4-4 for more context). In this project, drilling engineers sought to use a reinforcement learning algorithm to explore different calibration alternatives, and select the solution that best matched the well's characteristics in the shortest possible time. The algorithm imitates the mental processes of the expert who designs the fracturing treatments; it runs multiple scenarios until it finds a combination of parameters that match the well's pressure system. If the mini frac is poorly executed, it can cost ACME several millions; therefore the process was redesigned so an expert could review the calibration, approving or adjusting it according to their expertise. In this case, the algorithm did not replace completely the need for creativity. The team remained responsible for discovering the problem and defining goals and appropriateness criteria to measure the quality of the algorithmic solution.

Eliminating non-value-added activities and artefacts

A final practice observed in projects that transformed work by re-engineering the process was eliminating activities and artefacts that were not adding value to the product or service offered. Organisations realised that some activities and artefacts in their processes limited their ability to reap benefits from algorithms – even if such activities and artefacts were once deemed necessary. This practice involved questioning outdated assumptions and long-established rules that were bottlenecks in the process.

Consider the project for using algorithms on credit card onboarding in the business service area. The team decided to automate parts of this process using a combination of robotic process automation and computer vision. Credit card onboarding was a multistep process involving tasks such as credit checking, rating and issuing physical credit cards. When they began thinking about using algorithms to redesign the credit onboarding process, members of business services began questioning the need for using physical credit

cards. While physical cards had traditionally been deemed necessary for the process decades ago – as a token of identification and security for customer transactions – it also increased the lead time of the card onboarding process, making them 'lose business' – as some members put it. The business services area decided that a physical card was the main bottleneck in the process, and so decided to allow users to load a digital credit card via their phones. Eliminating the need for physical cards and automating the due diligence checking process enabled business services to onboard new customers almost within a day (rather than a month).

4.4.3 Adjusting the knowledge regime

There is a fourth main set of findings that we have already hinted at. When knowledge workers re-engineered the process – by removing non-value-added activities, adding steps for assuring quality, and controlling the sources of variation – they simultaneously adjusted the 'knowledge regime'. We based our construct of a knowledge regime on the works of Pachidi et al. (2021) and Howard-Grenville and Carlile (2006), and define it as the combination of social, material and discursive arrangements that govern the specific practices that actors use to learn and perform their work; the structural arrangements that inform which actors have authority over knowledge construction; and the valuation schemes that control what knowledge and methods matter in the construction of knowledge. When knowledge workers re-engineered the process, they also triggered power dynamics to adjust the knowledge regime. Below we provide a few examples.

Shifting structure and authority arrangements

When knowledge workers focused on redesigning the process through which the knowledge was transmitted, they ended up engaging in a dynamic to reorganise the structure and authority arrangements. In some cases, this reorganisation stemmed from

employing knowledge workers who had the tools, skills and time to map the process to create a product or service. This led to creating new structures, such as parallel organisations focused on scaling the application of algorithmic technologies. One example was IO, an organisation parallel to the business service function, created to provide '*automating as a service*' to the knowledge workers within business services. IO was composed of developers who offered three core services: advisory and consulting, end-to-end DevOps and federated automation solutions. IO worked in an agile way, in which the knowledge workers collaborated with developers to redesign processes around algorithmic solutions. As a result, business services morphed into a biparty rectangular organisation within ACME – which differed from the traditional matrix and heterarchical organisations. This structure also changed the arrangement of authority within the business services. The knowledge workers in the function lost jurisdiction to developers and reported not only to the head of the function but also to scrum masters. Decisions on where and how to use algorithmic technologies were distributed across both business services and IO.

After redesigning the working processes to include algorithmic technologies, other cases also experienced a shift in the structure and authority arrangements. In the case of group audit, auditors who automated the data collection and analysis for conducting some of their tests, find themselves needing to create a new team within audit, the data analytics team, with associated new roles to be able to provide support to the auditors on an on-demand basis. Similarly, engineers and petrophysicists in the drilling function described that as a result of reorganising the process to apply algorithmic technologies, the organisation created new interfaces. '*So prior [to the adoption of algorithms], in my particular case, I was in drilling… and knew the community quite well. We knew who the key stakeholders were and whom we had to engage with. There were not too many*

interfaces. Whereas now, a lot of us are having to find, okay, who are the new interfaces? Now we kind of engage with [the cybersecurity team], so, key architects and additional security people are involved in my work.' Thus, as they reorganised the work, they included new interfaces and structures that also altered who made decisions about knowledge work.

Altering schemes for evaluating new knowledge

When knowledge workers design new ways of working with algorithmic technologies, they also affect the mechanisms for learning and valuing knowledge. Before the implementation of algorithms, if a knowledge worker gave a judgement, the quality of that judgement was rarely questioned. It was socially accepted that they were experts whose credentials and past experiences granted them the legitimacy to make decisions. When algorithmic technology projects were launched, however, the subgroup of knowledge workers developing the solutions put significant efforts towards quantifying and testing the quality of the knowledge constructed by algorithms. One senior engineer working on projects Alpha 1 and Alpha 2 shared that such quantification emerged because knowledge workers using algorithmic technologies needed metrics to decide how and when to trust the algorithmic recommendations to avoid damaging the quality of the service offered. '*So, on the first half of the year, we are doing some hardware in loop testing, potentially testing on rigs you know, making sure we have some, some extensive tests ahead of our first deployment*' (Senior Engineer). The testing emphasised translating the quality of knowledge into metrics, such as the accuracy and the coefficient of determination of the algorithmic recommendations. While at the beginning this quantification was only applied to the knowledge produced by algorithms, eventually the knowledge workers in the function started thinking about their own recommendations,

in the same way, extending the quantification of the quality of knowledge from algorithms to engineers and petrophysicists too.

In addition to this shift to quantification of the quality of constructed knowledge, knowledge workers also altered their learning practices to accept intractable algorithmic recommendations. At the beginning of Alpha 1, for example, engineers were concerned about the effects of being unable to explain calibration recommendations made by the reinforcement learning algorithm. As time progressed, however, they focused less on understanding intractable algorithmic recommendations, and more on embedding boundary conditions in the algorithm to ensure that recommendations were safe:

> *'Working with a black box algorithm may be challenging for some and may be easy for others. And that will depend on the project… On Alpha 1, we will connect to the rig control system. So, it's going to be good but you also got to have some faith in it that it does the right thing. We've got sort of safety parameters around it. So, it would never go badly wrong, but you will never know why it's recommended a certain set of parameters over other certain parameters.' (Senior Engineer)*

Other knowledge workers also focused on developing tools to display the algorithmic recommendations and allowed knowledge workers to compare the algorithmic recommendations with their expected results.

> *'Now if we have the reinforcement learning algorithm that comes up with a solution that's calibrated, and you can see the two curves, an expert can say with one look, that's a good or that's a bad calibration, right? It doesn't matter what the parameters are. Again, there's an end for that, which they can be within. The expert says I can do better than that, or actually, this is good enough... So in that scenario, the fact that it's a black box it's actually less important because you can see the output.' (Senior Advisor)*

Because of the intractable recommendations, engineers lost their ability to learn from doing the job. In the case of Alpha 2, they could no longer build causal links between

parameters and pressure outputs. To overcome this problem, they dislocated learning in time and place. That is, they created simulated environments whereby engineers could train and play with different calibration parameters to see the effects on the fracture pressure.

> *'That enables you to prepare your mind. Just like how I go to bed and think about how the job looks like that particular well we did five months ago and, so expect this pressure to climb at this rate… Right. So [the simulator] can enable us to look for those patterns ahead of time. And you have a lot more confidence in that pressure match because we have automated that calibration piece.' (Senior Engineer)*

Updating the skill requirements

Finally, organisations that adjusted the knowledge regime, also enacted dynamics to update the skill requirements to do the work. For example, engineers in the drilling case argued that Alpha 2 empowered less experienced engineers to carry out fracturing treatments, thereby minimising service variability and improving process resilience and endurance. Before using algorithms, only a limited pool of experts felt comfortable with designing hydraulic treatments. One senior advisor commented that many junior engineers '*(..) are afraid to stick their neck out into the unknown*', which negatively affected the number of recruits for drilling roles, and so, the endurance of drilling operations. By implementing Alpha 2, the drilling organisation empowered junior engineers to do the fracturing treatments. '*The automation, really psychologically, is a crutch that [junior engineers] can hang on while they're still making the key decisions*' (Senior Advisor); '*it's an aid to them in the proper decision-making*' (Senior Engineer); '*it's basically enabling [junior engineers] to make the right decisions and not end up in a space where the well is messed up*' (Senior Petrophysicist).

Some organisations not only adjusted the skills required to conduct the job within the company, but they also influenced the knowledge workers' skills in the field. The drilling engineering department, which maintained strong ties with universities for research and recruitment purposes, motivated professors to update the curriculum of their engineering schools so graduates had data science skills straight out of university. '*When they finish their engineering degree to come to industry, they should have some knowledge of data science, machine learning techniques, artificial intelligence and coding*' (Senior Engineer). The rationale was that by updating the curriculum and the skills required to perform the work, the drilling organisation ensured that the engineers had the abilities to build and interact in successful projects involving algorithmic technologies.

4.5 Discussion

Drawing on our multiple-case study at ACME, we showed that ***knowing how*** to transform knowledge work is an ongoing accomplishment in which organisations call on interpretations, motivations and available means to construct new work practices. Particularly, we identified three distinct dynamics involved in knowing how to use algorithms to transform knowledge work. These included decision-making dynamics related to overcoming organisational inertia, transformational dynamics associated with redesigning work by altering specific tasks or re-engineering whole processes, and power dynamics linked to adjusting the structures that enable work and learning. To use our findings to explore the digital transformation of work, we now integrate our findings into a process model of how organisations experience, interpret and use algorithmic technologies to transform knowledge work. We then build on this process model to highlight our theoretical contributions, to outline managerial implications, and to explore under what circumstances technology use drives improvement in knowledge work.

4.5.1 An integrated process model of adoption of algorithmic technologies in knowledge work

How do organisations experience, interpret and use algorithmic technologies to transform knowledge work? We answer this question with an integrated process model of our findings, presented in Figure 4-2. Beginning at the left, the model depicts algorithmic technologies as an agentic entity that triggers organisational dynamics to overcome inertia. These dynamics incorporate recursive interpretations of algorithms, which are motivated by actors' ambitions and fears, and reinforced by past encounters with technology. Interpreting algorithms results in four types of logic in dealing with the adoption of algorithms – coping, challenging, differentiating and deflecting. These logics, in turn, orient the organisation towards certain knowledge work transformation pathways: one in which they alter the task by, for example, triaging and standardising tasks; and one in which they re-engineer the process by, for example, controlling the sources of process variation and eliminating non-value adding activities. As organisations transform their practices, they may engage in power dynamics to adjust the knowledge regime. This adjustment incorporates, intentionally or not, practices for altering the schemes for valuing knowledge, shifting structure and authority arrangements, and updating skill requirements. Unlike the process pathway, the task pathway may be enacted without further changes in the knowledge regime.

Our findings further suggest that successful algorithm projects create path dependency, in which knowledge workers redeploy the same automation model over time. By automation model, we mean the pattern of a transformation pathway (altering the task or re-engineering the process) and an algorithmic technology (e.g., machine learning). As many authors have suggested, path dependency results from existing investments in specific technologies and cognitive frames that favour established interpretations of technology use. Path dependency has had both positive and negative

implications for organisations. In some cases, path dependency allowed organisations to enjoy the cost advantages of creating economies of scale, as the cost of reusing algorithms in other knowledge work was significantly lower. Consider, for example, the case of the business services area. After initial success using robotic process automation, they used their current infrastructure as a scaffold to create and deploy circa 240 more robotic process automation solutions across the area. In other cases, path dependency had negative implications; it prevented organisations from realising the need for a different algorithmic technology or pathway for automating other knowledge work.

The project for creating risk reports for directors in the internal audit area illustrates the negative consequences of path dependency. In this project, internal auditors failed to incorporate algorithms for co-constructing audit finding reports as they disregarded shifting the way they transformed work, from altering specific tasks to re-engineering the whole process. When the project started, internal auditors chose to use algorithms to mechanise the task of creating summaries of audit reports, as they had focused on altering specific tasks in previous projects, such as in the algorithm project for finding inconsistencies in financial statements. This approach led to disappointing results. An analysis of the reports created by the natural language processing algorithm used revealed that algorithms were confused by variations in terminology, for example, in the multiple acronyms that ACME used, decontextualising audits' findings. Words such as 'cow', which could refer either to the animal or in ACME's lexicon, 'Control of Work', were not differentiated. The algorithm's inability to grasp the contextualised meaning of these words created confusing risk findings. Further variations that confused the algorithm stemmed from the differences in language used by engineers, who run operational audits, and accountants, who run financial audits. We found that controlling for some of these variations would have helped the algorithm to produce more accurate

reports. After several iterations, the audit team asked ACF to separate operational and non-operational audits to control some language variations between these two contexts. The new algorithm demonstrated improvement, suggesting that transforming the entire process to control for language and audit type variations may have resulted in a better construction of knowledge.

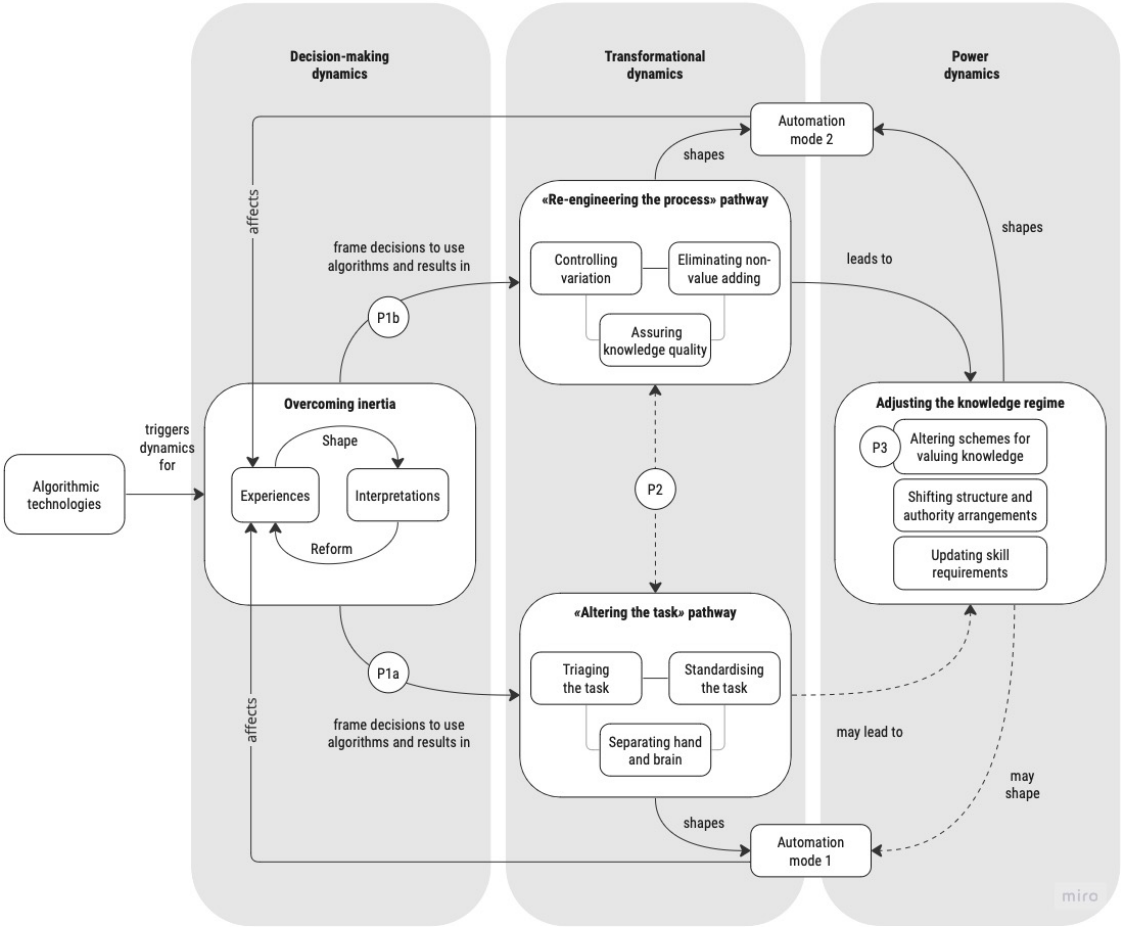


Figure 4-2 A process model of the adoption of algorithms in knowledge work

Adaptation and inertia in knowledge work

Our study contributes to our understanding of the adaptation to technological change and organisational inertia in knowledge work. Our findings concerned the cognitive

processes that shape the adoption of emergent technologies in knowledge work. In line with others (e.g. Eggers and Kaplan, 2009; Tripsas and Gavetti, 2000), we argued that cognitive representations and attention direct and restrict the search for solutions, shaping the adaptation to emergent technologies. Our study points, however, to internal dynamics in the transition to new technology-supported knowledge work, resonating with recent studies (e.g. Truelove and Kellogg, 2016). In particular, our grounded analysis shows social and decentralised decision-making processes at the core of these transitions. In our cases, we found that decisions on whether to adopt algorithms, in what kind of work, and to what extent depended on knowledge workers collectively experiencing and interpreting algorithms. These interpretations were motivated and socialised by actors with different statuses and positions in the organisation, not by dominant actors, as previous studies suggest. Socialising interpretations help knowledge workers construct cognitive frames, compelling distinct technology adoption logics – in our case, challenging, deflecting, coping and differentiating. Perhaps this social and decentralised decision-making is endemic to knowledge work, as the mastery of expertise grants knowledge workers more discretion than other types of work. Yet, as Barley and Kunda (2001) rightly argued, the nature of work has changed significantly across occupations, vesting discretion in less skilled workers and blurring the line between knowledge and other types of work. Thus, our findings might generalise to broader contexts. An avenue for future research is to investigate further how each actor's position in the organisation (formal and informal) shapes the adoption of technology and the pathways of transformation across different contexts.

4.5.2 Technological change in knowledge work

Our findings remedy the issue of the 'obscuring of variation' commonly present in previous studies on technological change (Bailey and Barley, 2020, p.2), in which

economists and policymakers had overlooked that work (and hence the tasks and required skills) differ across locales (c.f. Acemoglu and Autor, 2011; Arntz et al., 2016; Autor, 2015; Frey and Osborne, 2017). Undoubtedly, algorithmic technologies have given rise to a new wave of technological change and made some work susceptible to computerisation (Frey and Osborne, 2017). But as our findings show, job description databases, like the US Department of Labor's O*Net, do not accurately represent work in organisations, lessening their predictive power for understanding the effects of algorithms on labour. In line with other studies, our cases highlighted that work in organisations is complex, rich, localised, temporal, situated and co-constituted in practice (Brown and Duguid, 1991; Lave, 1988; Nicolini, 2011; Orlikowski, 2002). Take, for example, the task '*testing the design and operating effectiveness of controls*' performed by internal auditors. At the risk of oversimplifying our argument, even within the same organisation, the enactments of this task varied depending upon whether the task was embedded in compliance- or risk-based audits. Variations, in turn, had significant consequences for how technology changed internal audit work. When enacted in risk-based audits, the task had a dynamic and emergent flavour that hampered the use of algorithms. Thus, continuing to overlook the variation of work in our studies of technological change, would mean missing out on salient micro-dynamics that shape technological change.

Our findings have a second important implication for our understanding of technological change in organisations. Previous studies have suggested that three skills – creativity, social intelligence and physical manipulation – inhibit the use of algorithms for transforming knowledge work (Arntz et al., 2016; Frey and Osborne, 2017). In our field study, we found a somewhat more complex relationship between skills and the adoption of technology than that suggested. While complicating the use of algorithms, the need for these skills was not an inhibitor to technological change. Our cases report

the organisations' flexibility to deploy algorithms in work considered to require social intelligence and creativity. For example, drilling engineers used reinforcement learning algorithms for calibrating tailored fracturing and stimulation treatments – a task considered creative and subjective. Or, another example is workers in the business services area used RPA algorithms for credit card onboarding – a task that requires some social intelligence. In such projects, the work skills required acted as obstacles to using algorithmic technologies, but as with any obstacle, they can be removed or circumvented.

In addition to providing a corrective lens to assumptions of the homogeneity of work across organisations, and to the view that skills requirements determine the transformation of work, we used a knowing-in-practice lens for a third important purpose. Namely, to highlight that the transformation of knowledge work is a practice in and of itself, which organisations adapt, improvise and eventually learn through ongoing enactments. Our empirical findings highlight the tension between distinct enactments and path dependency. These findings support recent research that advocates a path-centric view on the adoption of technology and change (Pentland et al., 2022). By exploring how organisations have used algorithmic technologies to transform their work over time, we add detail to this literature. We found that when algorithmic technologies emerged, they created a new set of affordances that allowed organisations to follow multiple paths to do work – what Pentland et al. (2020) referred to as 'patterning'. For example, organisations that initially explored dynamics on the task pathway to activate algorithms, followed the same automation path in subsequent transformation projects. Shifting to a different pathway was hard, even when organisations failed to launch algorithmic technology projects with their preferred automation pathway. Thus, our study supports that a path-centric view of technological change is a valuable means for

studying technology transformation so that it can help understand the institutionalisation of technological change.

4.5.3 Technology use for improving knowledge work

Understanding work as situated, temporal and provisional enactments highlights the potential of algorithms in knowledge work. To date, the economic and learning perspective on technological change has attended to the negative consequences of algorithms, including the effects of labour displacement (c.f. Autor, 2015; Autor et al., 2003) and the dangers of using algorithms in skilled work (c.f. Balasubramanian et al., 2022; Lebovitz et al., 2021; Lebovitz et al., 2022). Our study challenges these findings, revealing that algorithms can not only be applied in knowledge work contexts but can also alleviate workloads and service pain points. The project on applying robotic process automation in credit card onboarding reduced service time, serving customers' needs better, while the project on using logic-based algorithms to build maintenance schedules let the maintenance scheduler cope with emergent demands for correcting asset faults without working extra time. Thus, continuing to overlook the potential of algorithms for creating improvements in efficiency and effectiveness, plus potential welfare aspects, would mean missing out on an important dimension of how algorithms shape the nature of knowledge work in organisations.

Why do algorithms drive improvements in some cases, and decrease the quality of outputs and learning in others? As we have hinted above, we argue that the heterogeneity of the consequences of algorithms for knowledge work depends on the shared practices for designing, embedding and using algorithms in established work. Our study identified a set of practices that were important enablers for successfully transforming knowledge work. These practices included, for example, triaging and standardising tasks and controlling for process variations; they allowed organisations to ensemble new work

practices in ways in which knowledge workers could, with algorithms, co-construct documents, services and judgements. In turn, such co-construction allowed more efficient knowledge work while maintaining quality standards. This finding resonates with recent studies on the role of *practices* through which knowledge workers deal with algorithm opacity (Anthony, 2021) and focuses on exploring how to redesign work to create synergistic effects between humans and algorithms (Sampson, 2021).

Drawing on the process model and discussion above, we now elicit some propositions on how organisations can use algorithms to improve knowledge work, starting with decision-making dynamics. It seems only logical that organisations beginning to transform their work decide to use algorithms to alter specific tasks, rather than engineering the entire process. After all, this is an appropriate approach to compartmentalise the risks of new technology and can help members to perceive technology as less threatening compared with bigger transformation alternatives (Truelove and Kellogg, 2016). However, such a strategy may be inappropriate in the long term. Because our analysis suggests that successful projects create path dependency, feeding existing cognitive frames and favouring existing transformation capabilities, organisations may fail to realise that a transformation pathway may be inappropriate for some technology change projects. This is the case, for example, when tasks that are tightly coupled or deal with high levels of variations in inputs, transformation and outputs. In those cases, a more holistic approach to work change is needed, one in which the entire process is considered and re-engineered. Failing to switch between pathways not only leads to wasting resources but also to missing opportunities for innovation. We, therefore, propose that:

Proposition 1: The potential for algorithms to improve knowledge work increases when the work transformation follows a process pathway instead of a task pathway.

We suggest that one way to overcome this path dependency and increase the likelihood of succeeding in using algorithms in knowledge work could be to sort and allocate potential transformation projects according to a system that considers variations and repetition in the process. We will refer to this practice as work triaging. Such an approach requires mapping the tasks involved in a process, noting inputs such as resources, dependencies and successors. The triage can also consider whether there are variations in the inputs (resources, the way they are received). By identifying variations, organisations may identify whether the simple substitution of humans for algorithms in a task is possible, or whether further changes in the process or inputs need to be considered before using algorithms in this task. Grounding this idea into the fit-to-context principle (Holweg et al., 2018), we propose:

Proposition 2: The potential for algorithms to improve knowledge work increases when organisations triage work to select the most suitable transformation pathway for the task at hand.

Based on the discussion above, we also suggest another essential practice that can increase the likelihood of improving knowledge work by using algorithms. This practice is *altering the schemes for evaluating the new knowledge*. We began this study by arguing that knowledge (or knowing) is inseparable from action; it is not fixed but provisional,

enacted and temporal (Orlikowski, 2002). The implication is that when an organisation changes its practices using algorithms, so does the 'knowing how' to do work and elicit recommendations. Thus, it would be unfair, unwise and even irresponsible to keep using old schemes for evaluating knowledge co-constructed by humans and algorithms. Of course, the challenge is that altering the schemes for evaluating knowledge triggers power dynamics in the organisation. In part, these power dynamics arise because knowledge workers protect the jurisdiction of expertise by preventing formalisations of knowing how to practice; in part, because knowledge workers prefer to sanction expertise by the demonstrated mastery of specific know-how practices (e.g.Abbott, 1988). Failing to realise that the new work practice, in which humans work with algorithms, requires new ways of evaluating the quality of produced knowledge outputs harms the organisation's performance. This does not mean that knowledge workers need to know how to fully comprehend the inner workings of algorithms (imagine scrutinising a transformer), but it means that knowledge workers need to revise how to measure, validate and monitor over time, and socialise understandings of the quality of analysis, recommendations and judgements co-constructed with algorithms. Recent studies support this view, suggesting that, for example, 'validating practices' have profound implications on the co-construction of knowledge (Anthony, 2021). We, thus, propose:

Proposition 3: The potential for algorithms to improve knowledge work increases when organisations evaluate and socialize the algorithm's knowledge outputs.

4.5.4 Managerial implications

Our findings have several implications for managers seeking to redesign knowledge work with algorithms. A key insight is that strategies involving the digital transformation of

knowledge work should be implemented considering that the ability to improve knowledge work with algorithmic technologies is primarily a function of the degree of novelty, rather than the required skills to perform a task as present studies would commonly suggest (Frey and Osborne, 2017). That is, to decide how to reallocate work between humans and algorithmic technologies, organisations should differentiate between work that involves well-defined and emergent problems. The former involves standard schemata and are good candidates for algorithms, such as '*auditing the financial statements of a company*'. The latter requires interpreting novel and unfamiliar events to create collective agreement on what is going on and plan potential solutions (Maitlis and Sonenshein, 2010; Weick, 1988, 1993), such as '*assessing the impact of COVID-19 on the productivity of the company*'. These tasks involving emergent problems make poor candidates for algorithms; humans should perform them using intuition, collaboration, imagination and abductive reasoning.

We also highlight that it is indeed possible to automate knowledge work that requires creativity and/or social interactions, provided that the degree of novelty remains moderate. But our findings suggest that success depends on the company's ability to first adapt the process architecture in which the task is embedded, including removing unnecessary interdependences, streamlining communications, and re-engineering the process to maintain quality; and second, to alter the structure and learning mechanisms of knowledge work, such as updating the skill requirements to perform work, dislocating learning in time and space, and reallocating authority between cross-boundary knowledge workers. We further argue that the problem of automating work requiring physical manipulation skills is an economic one. While algorithms can be embedded into robots to perform activities that require manual dexterity and/or work in confined spaces, they may be too costly or risk the safety of any nearby workers. Thus, it is more likely that

organisations will separate manual and mental labour to enable the use of algorithmic technologies in the latter.

Our results show that companies should avoid using algorithmic technologies for processes with low levels of standardisation. Principally, those processes that are tailored to customers' needs or that depend on experts' discretion tend to benefit less from automation algorithms. This observation also implies that organisations are more likely to succeed at using technology for knowledge work when they redesign it, standardising as many steps and outcomes as possible. Based on our analyses, we further note that organisations benefit from using algorithms in work related to organisation processes that are in high demand. For one, such decidedly repetitive processes allow the company to scale algorithmic solutions and compensate for the economic hurdles that may emerge from remapping the work of experts and professionals.

4.5.5 Limitations and further research

As with all case-based research, the generalisability of our findings is a major concern. We acknowledge that the cases presented in this study aim for naturalistic generalisability (Pratt et al., 2006), rather than statistical generalisability, and aim at building insights from unique, rather than common, cases (Yin, 2017). Care must thus be taken when transferring our findings to other organisations. One key limitation stems from the budget and time constraints that led us to focus on a specific part of the phenomenon of algorithms in knowledge work; we have focused only on the characteristics of knowledge work that make it more amenable to algorithmic technologies. We thus did not examine other potentially influential factors, such as status, experience, unionisation and organisation type, that may affect how knowledge workers use and interpret algorithmic technologies to construct their work. It may be that knowledge workers in, say, professional service firms construct narratives about algorithmic technologies in

strategic ways, protecting their agency and work, and ensuring control of key tasks; or it may be that experts' and professionals' aversion to algorithmic recommendations prevents them from realising the performance algorithmic technologies in the long run. Thus, two open questions that remain are how do knowledge workers with different characteristics construct, resist and react to organisational efforts to redesign knowledge work with algorithmic technologies? And how do these social processes end up moderating the performance gains of algorithmic technologies? We pose these for future research.

One more crucial area for future research involves investigating the impact of various types of algorithmic technologies on knowledge work improvement. In this study, we opted to use the term 'algorithmic technologies' instead of explicitly mentioning 'Machine Learning (ML)' or 'Robotic Process Automation (RPA)' in certain instances. This approach was taken to maintain a broader and more encompassing scope within the research. While we specifically presented ML and RPA as instances of algorithmic technologies, it is essential to note that our findings were aggregated, preventing a direct comparison and contrast to determine whether certain results are attributed to a particular algorithm type. Therefore, future research should delve into whether the specific algorithms employed by knowledge workers act as moderators in enhancing knowledge work.

To achieve a more nuanced understanding, researchers can explore how different algorithm types influence the efficiency, accuracy, and effectiveness of knowledge work. This could entail conducting comparative studies between various algorithms and their impacts on different aspects of knowledge work. Moreover, investigating the interaction between human expertise and algorithmic assistance could shed light on the optimal ways to integrate algorithmic technologies into knowledge-intensive tasks.

4.5.6 Outlook

In the wider context of the digitalisation of operations, the use of algorithmic technologies to improve knowledge work is likely to become more prevalent as technology advances further. Developing a deeper understanding of how such algorithmic technologies will alter, and improve, processes in general will thus be critical for the field of operations management going forward. In this paper, we have sought to extend our understanding by examining how algorithmic technologies are used to construct improvement opportunities in a knowledge work context. Our empirical study contributes to the literature by highlighting two generic pathways for adopting algorithmic technologies, either by automating a single task or by re-engineering the entire process. We highlight the role of the knowledge regime in the transformation process, thereby qualifying commonly held beliefs that tasks requiring creativity and social interactions cannot be automated. First and foremost, our study highlights that our understanding of how algorithmic technologies will enable the improvement of knowledge work, and the digital transformation of operations more generally, remains nascent. Most recent advances towards multi-modal or even general-purpose AI systems are likely to spark further adoptions of algorithmic technologies to contexts that are so far deemed not to be automatable, therein challenging our accepted wisdom even further.

References

Abbott, A.D. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press, Chicago.

Acemoglu, D. & Autor, D.H. (2011). 'Skills, tasks and technologies: Implications for employment and earnings', in: Card, D., Ashenfelter, O. (Eds), *Handbook of labor economics*. Elsevier, Amsterdam, pp. 1043-1171.

Acemoglu, D. & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy, 128* (6), 2188-2244.

Agrawal, A., Gans, J. & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press, Boston, MA.

Amabile, T.M. (1983). The social psychology of creativity: A componential conceptualization. *Journal of Personality and Social Psychology, 45* (2), 357-376.

Anthony, C. (2018). To question or accept? How status differences influence responses to new epistemic technologies in knowledge work. *Academy of Management Review, 43* (4), 661-679.

Anthony, C. (2021). When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative Science Quarterly, 66* (4), 1173-1212.

Armour, J. & Sako, M. (2020). AI-enabled business models in legal services: From traditional law firms to next-generation law companies? *Journal of Professions and Organization, 7* (1), 27-46.

Arntz, M., Gregory, T. & Zierahn, U. (2016). *The risk of automation for jobs in OECD countries: A comparative analysis*. Paris: OECD Publishing.

Autor, D. & Salomons, A. (2018). Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brooking papers on economic activity,* (Spring 2018), 1-63.

Autor, D.H. (2015). Why are there still so many jobs? The history and future of workplace automation. *The Journal of Economic Perspectives, 29* (3), 3-30.

Autor, D.H., Levy, F. & Murnane, R.J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics, 118* (4), 1279-1333.

Bailey, D.E. & Barley, S.R. (2020). Beyond design and use: How scholars should study intelligent technologies. *Information and Organization, 30* (2), 100286.

Bailey, D.E., Faraj, S., Hinds, P., von Krogh, G. & Leonardi, P.M. (2019). Special issue of organization science: Emerging technologies and organizing. *Organization Science, 30* (3), 642-646.

Bailey, D.E., Leonardi, P.M. & Chong, J. (2010). Minding the gaps: Understanding technology interdependence and coordination in knowledge work. *Organization Science, 21* (3), 713-730.

Balasubramanian, N., Ye, Y. & Xu, M. (2022). Substituting human decision-making with machine learning: Implications for organizational learning. *Academy of Management Review, 47* (3), 448-465.

Baldini, I., Castro, P., Chang, K., Cheng, P., Fink, S., Ishakian, V., Mitchell, N., Muthusamy, V., Rabbah, R. & Slominski, A. (2017). 'Serverless computing: Current trends and open problems', *Research advances in cloud computing*. Springer, pp. 1-20.

Barley, S.R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly, 31* (1), 78-108.

Barley, S.R. (1990). The alignment of technology and structure through roles and networks. *Administrative Science Quarterly, 35* (1), 61-103.

Barley, S.R. (1996). Technicians in the workplace: Ethnographic evidence for bringing work into organizational studies. *Administrative Science Quarterly, 41* (3), 404-441.

Barley, S.R. (2008). 'Coalface institutionalism', in: Oliver, C., Meyer, R.E., Greenwood, R., Lawrence, T.B. (Eds), *The Sage handbook of organizational institutionalism*. Sage, London, pp. 491-518.

Barley, S.R. & Kunda, G. (2001). Bringing work back in. *Organization Science, 12* (1), 76-95.

Bechky, B.A. (2006). Talking about machines, thick description, and knowledge work. *Organization Studies, 27* (12), 1757-1768.

Bechky, B.A. & O'Mahony, S. (2015). 'Leveraging comparative field data for theory generation', in: Kimberly D. Elsbach, Kramer, R.M. (Eds), *Handbook of qualitative organizational research*. Routledge, London, pp. 200-208.

Benbya, H., Pachidi, S. & Jarvenpaa, S. (2021). Special issue editorial: Artificial intelligence in organizations: Implications for information systems research. *Journal of the Association for Information Systems, 22* (2), 281-303.

Brandon-Jones, A., Lewis, M., Verma, R. & Walsman, M.C. (2016). Examining the characteristics and managerial challenges of professional services: An empirical study of management consultancy in the travel, tourism, and hospitality sector. *Journal of Operations Management, 42-43* (1), 9-24.

Brown, J.S. & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science, 2* (1), 40-57.

Brynjolfsson, E., Rock, D. & Syverson, C. (2019). 'Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics', in: Agrawal, A. (Ed), *The Economics of Artificial Intelligence*. University of Chicago Press, Chicago, pp. 23-60.

Charmaz, K. (2006). *Constructing grounded theory: A practical guide through qualitative analysis*. Sage, London.

Choudhury, P., Starr, E. & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal, 41* (8), 1381-1411.

Corbin, J.M. & Strauss, A. (2008). *Basics of qualitative research: Techniques and procedures for developing grounded theory*, 3rd ed. Sage, Los Angeles.

Davenport, T.H. (2015). 'Process management for knowledge work', in: vom Brocke, J., Rosemann, M. (Eds), *Handbook on Business Process Management*. Springer, Berlin, pp. 17-35.

De Rond, M. & Lok, J. (2016). Some things can never be unseen: The role of context in psychological injury at war. *Academy of Management Journal, 59* (6), 1965-1993.

Dellot, B., Mason, R. & Wallace-Stephens, F. (2019). *The four futures of work. Coping with uncertainty in an age of radical technologies*. London: RSA.

DeSanctis, G. & Poole, M.S. (1994). Capturing the complexity in advanced technology use: Adaptive structuration theory. *Organization Science, 5* (2), 121-147.

Dougherty, D. & Dunne, D.D. (2012). Digital science and knowledge boundaries in complex innovation. *Organization Science, 23* (5), 1467-1484.

Dreyfus, H.L. (2014). *Skillful coping: Essays on the phenomenology of everyday perception and action*. Oxford University Press, Oxford.

Eggers, J.P. & Kaplan, S. (2009). Cognition and renewal: Comparing CEO and organizational effects on incumbent adaptation to technical change. *Organization Science, 20* (2), 461-477.

Faraj, S. & Pachidi, S. (2021). Beyond Uberization: The co-constitution of technology and organizing. *Organization Theory, 2* (1), 1-14.

Faraj, S., Pachidi, S. & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization, 28* (1), 62-70.

Frey, C.B. & Osborne, M.A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114* (2017), 254-280.

Gioia, D.A., Corley, K.G. & Hamilton, A.L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods, 16* (1), 15-31.

Glaser, B.G. & Strauss, A.L. (2017). *The discovery of grounded theory: Strategies for qualitative research*. Routledge, London.

Grant, R.M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal, 17* (Winter Special Issue), 109-122.

Hadjimichael, D. & Tsoukas, H. (2019). Toward a better understanding of tacit knowledge in organizations: Taking stock and moving forward. *Academy of Management Annals, 13* (2), 672-703.

Harvey, J. (2016). Professional service supply chains. *Journal of Operations Management, 42-43* (1), 52-61.

Hayes, R.H. & Wheelwright, S.C. (1979a). The dynamics of process-product life cycles. *Harvard Business Review, 57* (2), 127-136.

Hayes, R.H. & Wheelwright, S.C. (1979b). Link manufacturing process and product life cycles. *Harvard Business Review, 57* (1), 133-140.

Hesse-Biber, S.N. & Leavy, P. (2011). *The practice of qualitative research*, 2nd ed. Sage, London.

Hollebeek, L.D., Sprott, D.E. & Brady, M.K. (2021). Rise of the machines? Customer engagement in automated service interactions. *Journal of Service Research, 24* (1), 3-8.

Holweg, M., Davies, J., De Meyer, A., Lawson, B. & Schmenner, R.W. (2018). *Process theory: The principles of operations management*. Oxford University Press, Oxford.

Howard-Grenville, J.A. & Carlile, P.R. (2006). The incompatibility of knowledge regimes: Consequences of the material world for cross-domain work. *European Journal of Information Systems, 15* (5), 473-485.

Hutchins, E. (1995). *Cognition in the wild*. MIT Press, Cambridge, MA.

Huysman, M. (2020). Information systems research on artificial intelligence and work: A commentary on "Robo-Apocalypse cancelled? Reframing the automation and future of work debate". *Journal of Information Technology, 35* (4), 307-309.

Kaplan, S. (2011). Strategy and PowerPoint: An inquiry into the epistemic culture and machinery of strategy making. *Organization Science, 22* (2), 320-346.

Kellogg, D.L. & Nie, W. (1995). A framework for strategic service management. *Journal of Operations Management, 13* (4), 323-337.

Kellogg, K.C., Orlikowski, W.J. & Yates, J. (2006). Life in the trading zone: Structuring coordination across boundaries in postbureaucratic organizations. *Organization Science, 17* (1), 22-44.

Kellogg, K.C., Valentine, M.A. & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals, 14* (1), 366-410.

Knorr Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Harvard University Press, Cambridge, MA.

Knorr Cetina, K. (2010). The epistemics of information: A consumption model. *Journal of Consumer Culture, 10* (2), 171-201.

Lacity, M., Willcocks, L.P. & Craig, A. (2016). Robotic process automation at Telefonica O2. *MIS Quarterly Executive, 15* (1), 21-35.

Lave, J. (1988). *Cognition in practice: Mind, mathematics and culture in everyday life*. Cambridge University Press, Cambridge, UK.

Lave, J. & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press, Cambridge, UK.

Lawrence, B., Zhang, J.J. & Heineke, J. (2016). A life-cycle perspective of professionalism in services. *Journal of Operations Management, 42* (1), 25-38.

Lebovitz, S., Levina, N. & Lifshitz-Assaf, H. (2021). Is AI ground truth really "true"? The dangers of training and evaluating AI tools based on experts' know-what. *Management Information Systems Quarterly, 45* (3), 1501-1526.

Lebovitz, S., Lifshitz-Assaf, H. & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science, 33* (1), 126-148.

Levinthal, D.A. & March, J.G. (1993). The myopia of learning. *Strategic Management Journal, 14* (S2), 95-112.

Lewis, M.A. & Brown, A.D. (2012). How different is professional service operations management? *Journal of Operations Management, 30* (1-2), 1-11.

Maitlis, S. & Sonenshein, S. (2010). Sensemaking in crisis and change: Inspiration and insights from Weick (1988). *Journal of Management Studies, 47* (3), 551-580.

Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P. & Dewhurst, M. (2017). *Harnessing automation for a future that works*. McKinsey Global Institute.

March, J.G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2* (1), 71-87.

McAfee, A. & Brynjolfsson, E. (2012). Big data: The management revolution. *Harvard Business Review, 90* (10), 60-68.

Miles, M.B. & Huberman, A.M. (1994). *Qualitative data analysis: An expanded sourcebook*, 2nd ed. sage, London.

Muro, M., Maxim, R. & Whiton, J. (2019). *Automation and artificial intelligence: How machines are affecting people and places*. Brookings.

Nicolini, D. (2011). Practice as the site of knowing: Insights from the field of telemedicine. *Organization Science, 22* (3), 602-620.

Orlikowski, W.J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization Science, 3* (3), 398-427.

Orlikowski, W.J. (2002). Knowing in practice: Enacting a collective capability in distributed organizing. *Organization Science, 13* (3), 249-273.

Orlikowski, W.J. & Barley, S.R. (2001). Technology and institutions: What can research on information technology and research on organizations learn from each other? *MIS Quarterly, 25* (2), 145-165.

Orr, J.E. (1996). *Talking about machines: An ethnography of a modern job*. ILR Press, London.

Østerlund, C. & Carlile, P. (2005). Relations in practice: Sorting through practice theories on knowledge sharing in complex organizations. *The Information Society, 21* (2), 91-107.

Pachidi, S., Berends, H., Faraj, S. & Huysman, M. (2021). Make way for the algorithms: Symbolic actions and change in a regime of knowing. *Organization Science, 32* (1), 18-41.

Pentland, B.T., Mahringer, C.A., Dittrich, K., Feldman, M.S. & Wolf, J.R. (2020). Process multiplicity and process dynamics: Weaving the space of possible paths. *Organization Theory, 1* (3)

Pentland, B.T., Yoo, Y., Recker, J. & Kim, I. (2022). From lock-in to transformation: A path-centric theory of emerging technology and organizing. *Organization Science, 33* (1), 194-211.

Polanyi, M. (1967). *The tacit dimension*. Routledge & Kegan Paul, London.

Pratt, M.G., Rockmann, K.W. & Kaufmann, J.B. (2006). Constructing professional identity: The role of work and identity learning cycles in the customization of identity among medical residents. *Academy of Management Journal, 49* (2), 235-262.

Ryle, G. (1949). *The concept of mind*. Hutchinson, London.

Sampson, S.E. (2021). A strategic framework for task automation in professional services. *Journal of Service Research, 24* (1), 122-140.

Sandberg, J. & Tsoukas, H. (2011). Grasping the logic of practice: Theorizing through practical rationality. *Academy of Management Review, 36* (2), 338-360.

Schmenner, R.W. (1986). How can service businesses survive and prosper. *Sloan Management Review, 27* (3), 21-32.

Silvestro, R., Fitzgerald, L., Johnston, R. & Voss, C. (1992). Towards a classification of service processes. *International Journal of Service Industry Management, 3* (3), 62-75.

Staats, B.R., Brunner, D.J. & Upton, D.M. (2011). Lean principles, learning, and knowledge work: Evidence from a software services provider. *Journal of Operations Management, 29* (5), 376-390.

Starbuck, W.H. (1992). Learning by knowledge-intensive firms. *Journal of Management Studies, 29* (6), 713-740.

Susskind, R.E. & Susskind, D. (2016). *The future of the professions: How technology will transform the work of human experts*. Oxford University Press, Oxford.

Syed, R., Suriadi, S., Adams, M., Bandara, W., Leemans, S.J.J., Ouyang, C., ter Hofstede, A.H.M., van de Weerd, I., Wynn, M.T. & Reijers, H.A. (2020). Robotic process automation: contemporary themes and challenges. *Computers in Industry, 115*, 103162.

Thompson, J.D. (1967). *Organizations in action: Social science bases of administrative theory*. McGraw-Hill, London.

Tripsas, M. & Gavetti, G. (2000). Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal, 21* (10-11), 1147-1161.

Truelove, E. & Kellogg, K.C. (2016). The radical flank effect and cross-occupational collaboration for technology development during a power shift. *Administrative Science Quarterly, 61* (4), 662-701.

Tsoukas, H. (1996). The firm as a distributed knowledge system: A constructionist approach. *Strategic Management Journal, 17* (S2), 11-25.

Tsoukas, H. (2005). *Complex knowledge: Studies in organizational epistemology*. Oxford University Press, Oxford.

Tsoukas, H. (2012). 'How should we understand tacit knowledge? A phenomenological view', *Handbook of organizational learning and knowledge management*. John Wiley & Sons, Inc, Hoboken, NJ, USA, pp. 453-476.

Waardenburg, L., Huysman, M. & Sergeeva, A.V. (2021). In the land of the blind, the one-eyed man is king: Knowledge brokerage in the age of learning algorithms. *Organization Science, 33* (1), 59-82.

Weick, K.E. (1988). Enacted sensemaking in crisis situations. *Journal of Management Studies, 25* (4), 305-317.

Weick, K.E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly, 38* (4), 628-652.

Wenger, E. (1999). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press, Cambridge, UK.

Wolcott, H.F. (1994). *Transforming qualitative data: Description, analysis, and interpretation*. Sage, London.

Woodward, J. (1965). *Industrial organization: Theory and practice*. Oxford University Press, London.

Yin, R.K. (2017). *Case study research and applications: Design and methods*. Sage publications, Los Angeles.

Appendix A. Interview protocol phase I

Part I. Knowledge worker information

1. Tell me about your role in three sentences.
2. How many years have you worked with the company?
3. How many people are you responsible for?

Part II. Characterisation of knowledge work tasks and perception of fit for algorithmic technologies

4. What are the five main tasks in your work? And how often do you do them?

From a-e, please rate each task on a 0–10-point scale according to the following dimensions. 0 means that the dimension (e.g., standardisation) is NOT present (e.g. the task is not standard at all), and 10 means the dimension is fully present (e.g., the task has been completely standardised).

a) How standard is this task? Please explain your answer.
b) How much social interaction is required to complete this task? Please explain your answer
c) How much creativity/ingenuity is required to complete this task? Please explain your answer.
d) How much physical manipulation/dexterity do you need to complete this task? Please explain your answer.
e) In your view, how suitable for automation is this task? And why?
f) What technologies do you currently use to support this task?

Part III. Perception of potential areas to automate

5. Where do you see, in general, the greatest potential for automation within your area? Why?
6. Specifically, what role do you see for RPA? For ML? in your area?
7. Is there anything you would like to add?

Appendix B. Semi-structured interview protocol phase II

Part I. Description of projects using algorithmic technologies

1. Tell me about project X
    - How did the project start?
    - What technology are you using? And why?
    - What was the aim of the project (e.g., knowledge work task)?
    - Who executes it?
    - Who is involved?
    - What is the status?

Part II. Algorithmic technologies and organising

2. How did you adopt algorithmic technologies?
    - What steps have you taken to use algorithmic technologies in your work?
    - What has worked?
    - What hasn't worked?
    - What changed?
    - What obstacles have you faced and how did you overcome them?

Part III. Algorithmic technologies and improvement

3. Tell me about a positive experience automating work with algorithms.
    - What was the project about?
    - What was the impacted task?
    - How did the project improve knowledge work?
    - What were the impacted areas (e.g., quality, delivery, cost, morale, safety)?
    - How did you measure success?

4. Tell me about a negative experience automating work with algorithms.
    - What was the project about?
    - What happened?
    - What was challenging about the project?

- What were the impacted areas (e.g., quality, delivery, cost, morale, safety)?
- What is the status of this project?

Appendix C. Data structure

First-order codes	Second-order practices	Overarching dimension
▪ Maintenance operations: Fearing algorithms would make maintenance workers redundant ▪ Business services: Seeing the appeal of algorithms to alleviate the area's pain points ▪ Drilling operations: embracing algorithms while recognising that other tools could get the same job done in more precise ways	Experiencing algorithms	Overcoming Inertia
▪ Internal audit: The biggest hurdle is that if I transfer my expertise to a brain, in two or three years I may be out of a job ▪ Drilling operations: Stakeholders interpreted algorithms as high risk because many data science projects haven't gone anywhere ▪ Internal audit: Auditors interpreted algorithms as a threat to their jurisdiction of expertise	Interpreting algorithms	
▪ Internal audit: Deciding to use algorithms in a tedious and repetitive task to make leaders content ▪ Business services: Using algorithms for credit card onboarding to cope with high service times ▪ Drilling operations: Adopting algorithms to differentiate themselves from other practitioners	Deciding to use algorithms	
▪ Maintenance operations: looking out for repeatable and predictable tasks to activate algorithms. ▪ Internal audit: Classifying whether tasks require social intelligence for information gathering or for dealing with emergent situations (informal)	Triaging tasks	Altering the task
▪ Internal audit: Constraining creative tasks to ensure audit testing is consistent across audits ▪ Deploying reinforcement learning to bridge the gap on the physics and art of fracturing on soft rock	Standardising tasks	
▪ Business services: Automating keystrokes to enter information in applications (transactional applications) ▪ Maintenance operations: Separating maintenance activities into planning and execution	Separating hand and brain	
▪ Drilling operations: Automating fracturing onshore is easier because the wells behave in a standard way ▪ Business services: Identifying standard and static activities in credit card onboarding ▪ Standardising work around drilling to deploy project Beta	Identifying and controlling the process variation	Re-engineering the processes
▪ Drilling operations: adding steps for experts to review the hydraulic treatment calibration suggested by an algorithm ▪ Business services: Creating dashboards for monitoring algorithms overtime.	Adding steps for assuring output quality	

▪ Business services: Assigning analysts to monitor the performance of RPA bots		
▪ Business services: Adjusting policies for using virtual credit cards ▪ Business services: Eliminating the need for issuing physical credit cards ▪ Drilling operations: Compressing the drilling process by incorporating the mini frac into the main hydraulic and stimulation treatment	Eliminating non-value-added activities and artefacts	
▪ Drilling operations: Dealing with new interfaces at work ▪ Internal audit: Engaging with the data analytics team to run audit tests ▪ Drilling operations: Reorganising teams based on member's expertise with the Bonsai platform	Shifting the structure and authority	
▪ Drilling operations: Using machine learning to look for patterns ahead of time and learn what was going on ▪ Drilling operations: Accepting that reinforcement learning does not need to provide explanations to provide value to the firm ▪ Business services: Testing models extensively before deploying	Altering practices for learning & valuing knowledge	Adjusting the knowledge regime
▪ Drilling operations: Using algorithms to empower engineers without specialised knowledge to run mini fracs. ▪ Business services: Automation improves the standardisation of a process and decreases the need for experienced workers	Updating the skill requirements	

Chapter 5

Enacting artificial intelligence: Connecting the development of machine learning systems with their operational effectiveness

Abstract

In this article, we propose a theory that connects the development of machine learning systems with their operational effectiveness. We do so by focusing on two key implications that impact the effectiveness of these systems: epistemic - those concerning the validation of knowledge – and normative - those pertaining to moral considerations. We argue that development practices shape the algorithm's operational effectiveness through three key mechanisms: cultivating a production culture, shaping production activities, and adopting production artefacts. We posit that while a production culture of continuous improvement can enhance an organisation's ability to address the epistemic implications of machine learning systems, it is less effective in tackling their normative implications. To effectively address normative implications, organisations need to standardise and reinforce production activities that focus on evaluating bias and inequalities present in historical data in the particular context of application, in addition to establishing a production culture of continuous improvement. We also suggest that production artefacts used to automate the creation of machine learning systems may pose short-term challenges but could ultimately reduce both epistemic and normative implications in the long run, given that the organisation has appropriately established a production culture and activities. By providing insights into the interplay between machine learning and operations, our theory offers practical guidance for scholars and practitioners engaged in the development and use of Artificial Intelligence. Our conceptual model promotes the responsible and effective integration of Artificial Intelligence systems into knowledge work.

***Keywords*:** MLOps, Machine Learning, Artificial Intelligence, Development, Practice.

5.1 Introduction

In March 2023, the Future of Life Institute published an open letter, endorsed by over 27,500 signatories, including industry leaders, urging for a halt in the development of Artificial Intelligence (AI). Notable examples of AI - computer programmes that rely on statistical inference to "learn" patterns in historical data without being explicitly programmed to imitate human decision-making – including machine learning systems and large language models. The letter voiced concerns about the "profound risks that AI poses to society and humanity", emphasizing the need for a deeper understanding of AI implications.

Concerns about AI systems have also emerged in academic circles. The literature acknowledges that while these systems hold the potential to foster innovation and enhance established processes (Chuang et al., 2021; Jung et al., 2018; Senoner et al., 2022), they also engender risks that could harm users and disrupt the organisation's operations. AI systems are characterized by high opacity, comprehensive digitization, anticipatory quantification, and hidden politics (Faraj et al., 2018), leading to two main implications. First, they pose epistemic, or knowledge-related, implications. AI-generated outputs are uncertain and probabilistic and their quality is difficult to scrutinize, complicating decision-making and hindering expertise development (Lebovitz et al., 2021). Second, they pose normative implications (Kane et al., 2021; Mittelstadt et al., 2016; Tsamados et al., 2021). These involve moral aspects such as biases, power asymmetries and eroded human agency (Anthony, 2021; Curchod et al., 2020; Kellogg et al., 2020). Both types of implications may arise from either organisational' actors inappropriate character or from inadequate capabilities to develop AI systems, and can harm the organisation's performance (Holweg et al., 2022). Studies on misconduct and failure (e.g. Ashforth et al., 2008; Collins et al., 2009; Greve et al., 2010) furnish valuable

insights that can be extended to understand how actors' character might shape the implications of AI. But the direct connection between an organisation's ability to develop AI systems and their operational effectiveness requires further exploration.

We seek to contribute to discussions about the relationship between AI and organisations by delving into the impact of AI development practices on operational effectiveness implications. Our focus is on an emerging set of practices known as "MLOps". Broadly defined, MLOps refers to principles, practices and tools used to design, deploy and maintain machine learning models (Alla and Adari, 2021; Gift and Deza, 2021; Treveil et al., 2020). Understanding the interplay between the development practices of machine learning systems and operational effectiveness is essential, as it underscores the two-way relationship between technology and organizing. That is, new technologies are not causal nor neutral, but they are an emergent resource shaped and reshaped by organisations' established practices (Bailey and Barley, 2020; Faraj and Pachidi, 2021; Leonardi, 2011). Furthermore, focusing on development practices, particularly MLOps, is essential due to the rising demand from "non-digital native" organisations to develop, deploy, and scale their machine learning systems (Benbya et al., 2020; Vartak, 2022). These organisations encounter unique challenges and potential unintended outcomes distinct from those associated with traditional software development (Faraj et al., 2018; Lebovitz et al., 2021; Mayer et al., 2020; Zhang et al., 2020). Thus, exploring the implications of AI development practices on operational performance remains a critical gap for advancing our understanding of this rapidly evolving field.

To explore these implications, we develop a theoretical framework that draws on three ideas from studies of technology in organisations, emphasizing the interplay between social and material factors in shaping human experience and organisational

practices (Cecez-Kecmanovic et al., 2014; Jones, 2014; Orlikowski, 2007, 2010; Orlikowski and Scott, 2008; Scott and Orlikowski, 2014). First, we adopt a *relational* perspective, underscoring the importance of the relationships between managers, developers and users in forming machine learning systems. Second, we embrace a *performativity* perspective, which recognises that organisational phenomena are dynamic, requiring 'specific and ongoing materializations' (Scott and Orlikowski, 2014, p.875). Finally, we consider the potential *agency* of material artefacts like data, scripts and models, and their role in shaping machine learning recommendations. This type of agency involves the ability of artefacts to act autonomously without human intervention, distinct from human agency which involves setting and achieving goals (Leonardi, 2011). By integrating these ideas, we can envision development practices as a central ground for configuring machine learning algorithms, taking into account both the social and material aspects at play.

This study contributes to studies on technology and operations by advancing our understanding of how AI development practices shape operational effectiveness consequences. We argue that these practices shape the operational effectiveness through three key mechanisms: cultivating a production culture, shaping production activities, and adopting production artefacts. Specifically, we suggest that the production culture affects an organisation's ability to address epistemic concerns by modifying the epistemic stance of the actors involved in the development of AI. Also, we further argue that production activities affect normative concerns by regulating the actors' behaviours that shape AI. Lastly, we argue that production artefacts affect both epistemic and normative concerns influencing the speed of AI development and operation. Broadly speaking, we advance the understanding of how algorithmic technologies and organisations co-constitute each other, providing actionable guidance for scholars and practitioners grappling with the use of algorithms in knowledge work (e.g. Anthony, 2021; Kellogg et

al., 2020; Waardenburg et al., 2021). These findings can help organisations design algorithmic technologies that advance knowledge work while mitigating adverse outcomes such as propagating inaccuracies, entrenching biases, or perpetuating inequalities.

This theory paper is structured as follows: in the first section, we build the historical and conceptual foundations for our discussion of MLOps. We begin by providing a brief overview of three precursors of MLOps that have defined how practitioners organise work in the context of software development: Agile, DevOps and machine learning. After exploring how MLOps emerged, we go on to integrate studies in engineering, data science and information systems to distil common themes in writing about MLOps. Then, in the second section, we situate MLOps in organisations by drawing on studies of technology in organisations. From this literature, we derive three important axioms that will guide the discussion of this study, namely the relational, agency and performativity axioms, and connect these axioms to the common themes in MLOps. In the third section, we use this theoretical apparatus to explore how MLOps shapes the epistemic and normative implications of algorithmic technologies in organisations. We conclude this manuscript by addressing the implications of MLOps for organisations, as well as identifying several theoretical and methodological opportunities for further research in this nascent area.

5.2 The emergence of the MLOps development practice

5.2.1 Historical foundations of MLOps

The roots of modern algorithmic technology development are linked to the advent of Agile software development and Lean methodologies. Although writings on Agile can be traced back to the 1990s and the Iacocca Institute (Burgess, 1994; Dove, 1991), it was

not until the beginning of the 2000s that Agile work began to emerge as a distinct topic of study, when practitioners translated the principles of Lean manufacturing into the context of software development (e.g. Fowler and Highsmith, 2001; Poppendieck and Poppendieck, 2003). In the more than 20 years since Agile entered the literature, practitioners have sought to replace the incumbent waterfall development approach, using Agile as a means to organise software development work to increase responsiveness to changing customer needs through faster development cycles, affording *'(…) the ability to thrive and reconfigure'* organisations (Krishnamurthy and Yauch, 2007, p.590) *'(…) to provide customer-driven products and services in a fast-changing market environment*' (Yusuf et al., 1999, p.37).

Various Agile methodologies emerged during the early 2000s, including Extreme Programming, Kanban and Scrum. These methodologies brought ideas of customer collaboration, feedback and adaptation that challenged the linear, sequential, and structured software developing approach promoted in 'Waterfall'. Extreme programming emphasized ideas of frequent feedback, continuous delivery, pair programming and test-driven development (Beck, 1999; Beck and Fowler, 2001); Kanban focused on visualizing the workflow, limiting work in progress and optimizing the flow of value (Ahmad et al., 2018; Ahmad et al., 2013); while Scrum paid attention to defining roles, events and artefacts for iterative and incremental product delivery (Schwaber, 1997, 2004). These methodologies shared one common goal: to help organisations deal with changing requirements – an ability that was mostly non-existent in Waterfall development, which suffers from an overreliance on following the original plan and meeting the original customer requirements (Boehm, 1988; Royce, 1987). Evidently, such plans and requirements were most of the time inaccurate, caused by information asymmetries, and unable to deal with emergent needs. Recognising the importance of responding to change,

Agile methodologies surfaced as value-based cyclical development alternatives that aimed at incorporating customer feedback continuously and rapidly.

By the 2010s, however, practitioners expressed discontent towards Agile as it failed to achieve many of the originally intended outcomes. Humble and Farley (2010) argued that Agile failed to drive wide performance improvements because it mostly focused on development – work related to planning, development, and testing of software – and neglected operations – work related to software delivery, infrastructure management, security and compliance assurance, and maintenance. This resulted in organisations overlooking critical activities that ensure customers can access and use algorithmic-based solutions appropriately, leading to various issues. Integration of code with the broader organisation's infrastructure caused longer development cycles, increased costs, and quality issues (Humble and Molesky, 2011). The substantial differences between development and production environments made some pieces of code obsolete, resulting in delays in software releases and straining organisational operations. In summary, Agile's focus on development activities at the expense of operations led to a failure to achieve intended outcomes, causing challenges for organisations that relied solely on Agile.

In response to these challenges, the practice of "DevOps" emerged as a solution aimed at reducing the software development cycle by breaking silos between "Development" and "Operations" work. DevOps integrates testing, deployment and release activities into the development process (Dyck et al., 2015; Humble and Farley, 2010; Humble and Molesky, 2011), and draws inspiration from Operations Management, specifically the theory of constraints advanced by Goldratt and Cox in "The Goal" (1984) and "Critical Chain" (1997), DevOps aims to identify and remove constraints, or "bottlenecks", in the production of algorithm-based solutions. DevOps has gained popularity, fuelled by influential works such as "The Phoenix Project" (Kim et al., 2014),

which has become a seminal reference in the field. Also, leading technology companies like Amazon have contributed to its diffusion by creating knowledge bases, infrastructure and services to simplify and scale its adoption. However, while DevOps has become a popular approach to developing algorithm-based solutions, interest in it has primarily remained within practice communities, with most references found in books, social media, and conference articles (Leite et al., 2019; Mishra and Otaiwi, 2020).

While the emergence of DevOps transformed software development practices and services, a pivotal technological disruption emerged in the 2010s with the advent of machine learning. Unlike traditional software development, machine learning does not require developers to specify instructions for achieving a particular outcome, but rather relies on statistical inference to identify patterns in historical data and learn those instructions. Despite machine learning having a long history, with models dating back over 80 years that imitate human neural networks and learning (e.g. McCulloch and Pitts, 1943; Rumelhart and McClelland, 1986), it was not until 2006 that the approach began to gain popularity among non-digital native companies. This newfound interest was catalysed by the works of Geoffrey Hinton, Youhua Bengio, and Yann LeCun, who advanced a machine learning architecture called "Deep Learning" (e.g. Bengio et al., 2006; Bengio and LeCun, 2007; Hinton et al., 2006), following the Neural Computation and Adaptive Perception (NCAP) research initiative. For a comprehensive history of deep learning, refer to Goodfellow et al. (2016, pp.12-26).

In recent years, the field of machine learning has undergone significant advancements in both scale and scope, enabling the creation of machine learning models that emulate human expertise across various domains, including healthcare (e.g. Lebovitz et al., 2021; Lou and Wu, 2021), law (e.g. Armour and Sako, 2020; Zhang et al., 2020), and finance (e.g. Cao, 2022; Mayer et al., 2020). From identifying people struggling with

suicide based on online posts (Chau et al., 2020), to recommending improvement actions to improve production quality (Senoner et al., 2022), Machine learning models are capable of identifying patterns otherwise missed by experts (due to complexity or scale), thereby, providing organisations with a means of fast value creation. The availability of large datasets, increased computational power, and attention-based architectures (Vaswani et al., 2017), have further driven the development of machine learning. For example, large language models such as ChatGPT, DALL-E, and Bard, have emerged, prompting even non-digital native organisations to rethink their value proposition through the integration of machine learning-based solutions. However, it is worth noting that, despite the excitement around ground-breaking deep learning, most practical adoptions of machine learning among non-digital native organisations to date continue to rely on classical machine learning algorithms, such as classification and regression.

As organisations leverage machine learning to redesign their products, processes, and services, they are increasingly aware of the unique challenges that accompany them (Faraj et al., 2018; Lebovitz et al., 2021; Martin, 2019; Mayer et al., 2020; Sculley et al., 2014). While machine learning can accelerate the development of products and services, it also presents new issues, including hidden values embedded in algorithm recommendations (Faraj et al., 2018), inaccurate predictions due to uncertain labels (Lebovitz et al., 2021), and lack of explainability (Hafermalz and Huysman, 2022).

A growing interest in applying DevOps principles to ensure efficient and effective delivery of machine learning has emerged. This has given rise to a new practice, known as "MLOps", which aims to address the unique challenges of deploying and maintaining machine learning models in production environments. In the following, we synthesize existing literature on MLOps towards defining the core concepts and exploring key themes.

5.2.2 Conceptual foundations of MLOps

Although the idea of MLOps has started to spread among practitioners and scholars alike, considerable variation remains in how MLOps is understood. Kreuzberger et al. (2022) talk of MLOps as a 'paradigm', describing it as a set of best practices, constructs, and development culture. Some refer to MLOps practices (Hewage and Meedeniya, 2022), while others write of MLOps as an application of DevOps to machine learning (Alla and Adari, 2021; Renggli et al., 2021). In this section, we discuss the common themes in writing about MLOps.

MLOps is often used without a clear definition, and even when defined, it is used with a wide range of meanings. The differences in definitions reflect salient underlying ontological assumptions about what MLOps entails. One key ontological difference reflected in various definitions is, for example, whether MLOps refers to a collection of tools and techniques or to such collections in use - i.e., what organisational actors do with these tools in their situated action when dealing with boundaries. Some definitions frame MLOps as a toolkit focused on the development and deployment of machine learning models (Symeonidis et al., 2022). This view is echoed in other conceptualizations of MLOps that describe it as the '[…] assortment of tools … [for] implementing MLOps principles" (Alla and Adari, 2021).

Other definitions position MLOps as a collection of tools *in use* that is enabled by organisational actors' intentionality, with the effects of MLOps in the organisation being co-concreated by both human and material agencies. This view of MLOps is often described as a 'practice' or a 'paradigm'. For instance, Granlund et al. (2021) noted the notion of human intentionality in MLOps by defining it as '[…] *advocating* automation and monitoring at all steps of the [machine learning] system' (p.83). Similarly, Kreuzberger et al. (2022) describe how the effects of MLOps are guided not only by the

tools themselves but also by the development culture, the knowledge of organisational actors, and the working practices. In line with this perspective, our focus in this article is on the use of MLOps within organisations rather than simply on the tools themselves.

Table 5-1 Selected definitions of MLOps

Author	**Themes**	**Nature of definition**	**Definition**
Alla and Adari (2021)	Principles	Application of DevOps	"[T]he intersection between machine learning and DevOps practices. (…) MLOps adopts DevOps principles and applies them to machine learning models in place of software, uniting the developmental cycles followed by data scientists and machine learning engineers with that of operational teams to help ensure continuous delivery of high-performance machine learning models"
Granlund et al. (2021)	Activities	Process	"MLOps – amalgamation of ML and operations – refers to advocating automation and monitoring at all steps of ML system development and deployment. As with any piece of software, support is needed for system integration, testing, releasing, deployment, monitoring and infrastructure management."
Hewage and Meedeniya (2022)	Activities and Principles	Life cycle	"The MLOps practice brings ML models into the software solution production process. It bridges the ML applications with DevOps principles, where deployments and maintenance of ML models can be automated in the production environment (…). MLOps systems should be capable of acting as collaborative, continuous, reproducible, tested, and monitored systems to achieve organisational MLOps goals. The development *life cycle* of MLOps consists of three major components as data, model, and code."
Kreuzberger et al. (2022)	Activities and Principles	Paradigm	"MLOps (Machine Learning Operations) is a paradigm, including aspects like best practices, sets of concepts, as well as a development culture when it comes to the end-to-end conceptualization, implementation, monitoring, deployment, and scalability of machine learning products. Most of all, it is an engineering practice that leverages three contributing disciplines: machine learning, software engineering (especially DevOps), and data engineering. MLOps is aimed at productionizing machine learning systems by bridging the gap between development (Dev) and operations (Ops). Essentially, MLOps aims to facilitate the creation of machine learning products by leveraging these principles: CI/CD automation, workflow orchestration, reproducibility; versioning of data, model, and code; collaboration; continuous ML training and evaluation; ML metadata tracking and logging; continuous monitoring; and feedback loops."

Author	Themes	Nature of definition	Definition
Mäkinen et al. (2021)	Activities	Practice	“MLOps refers advocating automation and monitoring at all steps of ML system development and deployment, including integration, testing, releasing, deployment and infrastructure management.”
Muralidhar et al. (2021)	Activities	Body of work	“MLOps—Machine Learning Operations—refers to the body of work that focuses on the full lifecycle of ML model deployment, performance tracking, and ensuring stability in production pipelines”
Raj et al. (2021)	Tools and artefacts	Toolkit	“MLOps is presented as a solution for handling and improving AutoML and feature-specific ML engineering, such as explainability and interpretability.”
Renggli et al. (2021)	Principles	Application of DevOps	“The term “MLOps” is used when this DevOps process is specifically applied to ML [4].”
Ruf et al. (2021)	Principles	Discipline	“Formally, MLOps is a discipline which is formed of a combination of ML, Development Operations (DevOps), and Data Engineering to deploy ML systems reliably and efficiently”
Symeonidis et al. (2022)	Artefacts	Toolkit	“MLOps (machine learning operations) stands for the collection of techniques and tools for the deployment of ML models in production [12]. Contains the combination of DevOps and Machine Learning. (…) MLOps attempts to automate Machine Learning processes using DevOps practices and approaches.”
Tamburri (2020)	Artefacts	Orchestration of tools	“In essence, Machine Learning Operations—or MLOps—entail the complex cloud the collection of techniques and tools for the deployment of ML models in production realising at least 5 functions: (1) data ingestion/ transport; (2) data transformation; (3) continuous ML (re-) training; (4) continuous ML (re-)deployment; (5) output production/presentation to the end-user (e.g., in the form of business intelligence [4]).”
(Testi et al., 2022)	Activities	Process	"MLOps can be viewed as the iterative process of pushing the latest best ML models to production"
Google	Principles Activities	Paradigm	“MLOps is an ML engineering culture and practice that aims at unifying ML system development (Dev) and ML system operation (Ops). Practicing MLOps means that you advocate for automation and monitoring at all steps of ML system construction, including integration, testing, releasing, deployment and infrastructure management.”

Common themes in writings on MLOps

Despite these notable differences, there are also several recurrent themes across definitions of MLOps. First, *principles* play a key role in MLOps. This view is anchored in the notion of MLOps as a 'development culture' (Kreuzberger et al., 2022) or as an 'engineering culture' (Google, 2020; Leonardi and Barley, 2008) shaped by a set of principles that guide '(…) how things should be realized …' (Kreuzberger et al., 2022). This is captured in descriptions of MLOps as the integration of software engineering principles (Zhou et al., 2020), DevOps principles (Hewage and Meedeniya, 2022; Raj et al., 2021; Renggli et al., 2021) and data design principles (Tamburri, 2020) into machine learning development. There seems to be a lack of agreement on what the principles of MLOps are. Particularly, some authors conflate principles with activities. We take principles to mean standards that govern developers' behaviour. After analysing the literature, we consolidate these principles into three: continuous improvement, collaboration and reproducibility. Some authors suggest versioning as a principle instead of reproducibility, but we view them as interchangeable because they both aim to ensure that results can be replicated and verified by others. For descriptions of these principles, please see Table 5.2.

Although studies elucidate the principles of collaboration and reproducibility with relative clarity, the definition of 'continuous improvement' in MLOps tends to lack specificity. While concepts such as 'iterative-incremental processes' (Visengeriyeva et al.) and activities like continuous deployment, continuous training, and continuous integration encompass the idea of continuous improvement, a precise definition is frequently absent. Even when a definition is provided, continuous improvement is commonly used vaguely as a guiding principle in MLOps, emphasizing the importance of consistently striving for improvement in the machine learning development process. This observation is not surprising, considering that even in the Operations Management

literature, where the term originated, continuous improvement has been subject to diverse interpretations.

To address this lack of clarity, we aim to clarify the meaning of continuous improvement in the context of machine learning. Continuous improvement or 'kaizen' has been defined in various ways. Some scholars focus on the outcomes, highlighting 'improvement initiatives that increase successes and reduce failures' (Juergensen, 2000). Others emphasize an ongoing process, describing it as *'a company-wide process of enabling a continuing stream of focused incremental innovation'* (Bessant et al., 1994, p.17). Additionally, some scholars underscore the importance of culture and behaviour, characterizing it as '*a culture of sustained improvement targeting the elimination of waste in all systems and processes of an organisation'* (Bhuiyan and Baghel, 2005, p.761).

While these definitions offer valuable insights, they also possess limitations and ambiguities. For instance, the first definition may lead to the misconception that continuous improvement consists of one-time initiatives rather than emphasizing its continuous nature. The second definition raises questions about whether continuous improvement must be limited to company-wide efforts, overlooking the potential for individual teams to engage in continuous improvement. Although the last definition acknowledges the role of culture, it fails to specify the specific type of culture that fosters sustained improvement. Furthermore, it neglects to recognize that continuous improvement is fundamentally a change project, as emphasized by its Japanese origins (kai - change, zen - good).

Addressing the aforementioned limitations, we adopt and refine the definition of continuous improvement based on Holweg et al. (2018, p.168), as:

'*Worker-engaged* (and usually led) activities that, *on a continuous basis*, review and improve the actual process *in alignment with well-defined goals*, through activities that promote quality circles, waste elimination and the like'.

This refined definition highlights three fundamental constituents of continuous improvement. Firstly, it acknowledges the necessity of sustained activity, emphasizing that continuous improvement entails ongoing efforts rather than being a one-time endeavour. Secondly, it recognizes the importance of involving frontline workers, appreciating their expertise and understanding that they are often best positioned to identify opportunities for improving their work processes (Ohno and Bodek, 1988; Sugimori et al., 1977). Lastly, it underscores the significance of establishing clear goals accompanied by continual evaluation. This aspect reflects the understanding that continuous improvement is guided by a shared objective, and progress is measured against these objectives.

Accordingly, within the context of designing and operating machine learning systems (MLOps), the principle of continuous improvement can be contextualized as: *Promoting data scientists-engaged (and usually led) activities that, on a continuous basis, review and enhance the machine learning system in alignment with well-defined goals, such as improving accuracy. These activities involve utilizing practices such as quality circles, waste elimination, continuous training, continuous deployment, continuous integration, and the like.*

Table 5-2 Examples of specific MLOps-related principles

MLOps -related principles	Definition
Continuous improvement	Promoting data scientists-engaged (and usually led) activities that, on a continuous basis, review and enhance the machine learning system in alignment with well-defined goals, such as improving accuracy. These activities involve utilizing practices such as quality circles, waste elimination, continuous training, continuous deployment, continuous integration, and the like
Collaboration	Foster collaboration between data scientists, machine learning engineers, and operations teams to facilitate communication and knowledge sharing.
Reproducibility	Refers to the ability to recreate or reproduce the results of a machine learning workflow, including the data processing, model training, and evaluation steps. This means that if the same data and code are used, the results should be identical or very similar, regardless of the environment or computing infrastructure.

Second, MLOps concerns a process, a sequence of *activities* to scale a data science project into production. This is captured in descriptions of MLOps as an approach to deal with the '(…) machine learning development lifecycle' (Garg et al., 2021), to '(…) advocate for automation and monitoring at all steps of ML system construction' (Google, 2020), and to make '(…) ML a standard operational procedure, which requires automation in the form of pipelines' (Mäkinen et al., 2021). Here, we should note that 'pipelines' refer to a sequence of processing procedures and routines for manipulating arranged so the output of each procedure serves as the input for the next (Alla and Adari, 2021). Other authors make references to the MLOps process as a 'workflow' comprised of 'phases', such as 'Designing the ML-powered application', 'ML Experimentation and Development', and 'ML Operations' (Visengeriyeva et al.). For examples and descriptions of MLOps-related activities, see Table 5-3.

Table 5-3 Examples of specific MLOps-related activities

MLOps-related activities	Definition
Data extraction	"(…) select and integrate the relevant data from various data sources for the [machine learning] task".(Google, 2020)

MLOps-related activities	Definition
Feature engineering	"Then feature engineering techniques are applied to extract and identify vital informative features for the design of the ML models (Hewage and Meedeniya, 2022, p.3) 16
Algorithm selection	"Hyper-parameter tuning and optimization processes are implemented before the training process."
Model development	"concerns experimenting and developing a robust and reproducible model training procedure (training pipeline code), which consists of multiple tasks from data preparation and transformation to model training and evaluation. (Salama et al., 2021, p.7)
Model testing	Unit testing and integration testing stages will be accompanied by the code build stage using a pipeline. Model testing and validation are also important to check the performance of the model. When the model complies with the expected accuracy level, it is deployed to the production environment.
Continuous delivery	"Continuous Delivery is a practice in which the software is built in a manner that is always in a production-ready state. This ensures that changes could be released on demand quickly and safely. The goal of CD is to get the new features developed to the end user as soon as possible." (Ruf et al., 2021, p.7)
Continuous integration	"Continuous Integration is a software practice that focuses on automating the process of code integration from multiple developers. In this practice, the contributors are encouraged to merge their code into the main repository more frequently. This enables shorter development cycles and improves quality, as flaws are identified very early in the process." (Ruf et al., 2021, p.7)
Continuous deployment	"Continuous deployment is a practice in which every change is deployed in production automatically." (Ruf et al., 2021, p.7)
Continuous training	Continuous training "(…) aims to automatically retrain the model where needed."
Continuous Monitoring	"Is about serving the model that is deployed in production for inference" (Salama et al., 2021, p.8)

Third, MLOps concerns a set of tools or artefacts that enable organisations to create and maintain machine learning models. This theme is captured in descriptions of MLOps as a '*collection of techniques and tools*' (Symeonidis et al., 2022; Tamburri, 2020), or MLOps as comprising *'(…) major components as data, model, and code*' (Hewage and Meedeniya, 2022), To support the development and maintenance of machine learning models, MLOps relies on various artefacts, such as data repositories, selected machine learning models, model registries, pipelines, and automated testing scripts. These artefacts play different roles in the machine learning development lifecycle, and they help

ensure that the models are developed, deployed, and maintained in a systematic and reproducible way. For more information about these artefacts, please refer to Table 5-4.

Table 5-4 Examples of specific MLOps-related tools and artefacts

MLOps-related tools and artefacts	**Definition**
Data repositories	A data repository is a central location or database where data is stored, managed, and accessed. It is typically used to organize and store large amounts of data in a structured and accessible way, making it easier to retrieve and analyse. The data contained is used for training and testing machine learning models.
Selected machine learning model	A machine learning model is a program or algorithm that learns from data to recognize certain patterns or relationships. It is a mathematical representation of a system or process that is trained on a dataset, and once it has learned from that data, it can be used to make predictions or recommendations on new, previously unseen data. The goal of training a machine learning model is to enable it to generalize to new situations beyond the training data. The choice of model and training process depends on the nature of the problem and the available data.
Model registry	"The model registry stores centrally the trained machine learning models together with their metadata. It has two main functionalities: storing the machine learning artifact and storing the machine learning metadata" (Kreuzberger et al., 2022, original abbreviates machine learning as ML)
Pipelines	"A pipeline is an infrastructure that contains a sequence of components manipulating information as it passes through the pipeline. The function of the pipeline can slightly differ within the setups, so be sure to refer to the graphs and explanations to get a better idea of how the pipeline in the example functions" (Alla and Adari, 2021) Some common pipelines include data pipelines and automatic training pipelines.
Automated testing script	Automated testing scripts are pre-written code designed to test software applications automatically, without the need for manual intervention. They can help to identify bugs, errors, and performance issues in a more efficient and reliable manner. These scripts also continue the metrics used to evaluate the performance of machine learning models.

Based on these common themes, we thus define MLOps as:

a development and operations practice that involves adopting principles of continuous improvement and collaboration, leveraging patterns in historical data and autonomous algorithms, and enacting a set of activities for building, testing and updating products or services iteratively.

5.3 Situating MLOps in organisations: A sociomaterial ontology

5.3.1 Sociomateriality

Thus far, we have focused largely on the historical and conceptual foundations of MLOps, highlighting three common themes in MLOps – principles, activities and artefacts. Before associating these themes with the consequences of machine learning on organisational life, we now turn our attention more fully to sociomateriality as an ontology of technology in organisations.

The concept of sociomateriality emphasizes the interplay between social and material factors in shaping human experience and organisational practices. The term suggests that material objects, such as technologies and physical artefacts, are not neutral and passive but are instead shaped by social and cultural factors, and in turn, shape social and cultural practices. As such, sociomaterial present a useful ontology for exploring technological phenomena in organisations. For example, sociomateriality provided a powerful lens for analysing how the introduction of robotic telepresence reshapes the coordination of medical work (Beane and Orlikowski, 2015). Similarly, Orlikowski and Scott (2015a) applied sociomateriality to explain how crowds and algorithms reshape anonymity in the tourism sector.

Although research on the relationship between the social and material has become significant in management and organisational theory (e.g. Barley, 1986; DeSanctis and Poole, 1994; Scott and Orlikowski, 2014), there is no unified theory of sociomateriality (Cecez-Kecmanovic et al., 2014; Jones, 2014). However, different sociomateriality theories can be categorised into substantialist ontologies that draw boundaries between organisational actors and artefacts (e.g. Leonardi, 2013; Leonardi et al., 2012), and relational ontologies that assume '*the social and the material and inherently inseparable*' (Orlikowski and Scott, 2008, p.456). In this study, we follow the latter and emphasise

three axioms: the relational, the agency and the performativity axiom. These axioms will be explored in greater detail in the following.

The relational axiom: Human experience (phenomenon) is shaped by entangled social and material factors. According to the relational axiom, the world is not composed of independent, self-contained, pre-existing entities, but rather, the world is made up of a constitutive entanglement between humans, organisations and materiality (Barad, 2003, 2007). In other words, everything that is social also has a material aspect to it, and everything that is material also has a social aspect to it (Orlikowski, 2007). This axiom presumes phenomenon is the basic unit of reality and relations are more fundamental than entities (Emirbayer, 1997). That is, individuals do not pre-exist in their interactions but rather emerge through and as part of their entangled intra-relating (Scott and Orlikowski, 2014). It calls for breaking down the "analytical boundaries between humans and technology" (Orlikowski and Scott, 2015b, p.455), and instead, views the world as a dynamic web of heterogeneous human and nonhuman actors that are co-constructed in practice (Latour, 2007). This perspective is not entirely intuitive due to semantic reasons (Cecez-Kecmanovic et al., 2014). But, it highlights that organisational practices 'always entail some sort of technological mediation' (Orlikowski and Scott, 2008, p.454). Moreover, since the social and material worlds lack independent existence in the absence of the other (Barad, 2003), neither realm can be said to hold a privileged status in organisational life. This challenges previous research and theory that focused on how social and material entities interact with each other (Orlikowski, 2010; Slife, 2004), and implies that technology does not have inherent properties, but rather emergent properties that arise from its relationships with other actors.

The performativity axiom: The world is a series of ongoing processes of materialization and meaning-making constantly co-constructed by human

and non-human agents. The performativity axiom suggests that the world is not a fixed set of objects and entities that exists 'out there', but it is instead a series of ongoing processes of materialization and meaning-making that are constantly co-constructed by human and non-human agents. That means that identities (Costas and Fleming, 2009), spaces (Lawrence and Dover, 2015), and markets (Callon and Muniesa, 2005) are performed and made real through our actions and interpretations. Callon and Muniesa (2005), for example, show that economic theories and models are not simply representations of reality, but they are *performative;* they actively shape and construct the economic phenomena they seek to explain. This notion of performativity, thus, challenges the idea that empirical reality is independent, prior, single, and determinate and instead offers a contingent, dynamic, and emergent understanding of reality as a practical accomplishment (Scott and Orlikowski, 2014). This requires a rethinking of the relationship between discursive practices and material phenomena. Accordingly, this takes on a post-humanist position where discursive practices are not solely human-based activities, but rather involve specific material configurations that enact local determinations of boundaries, properties, and meanings in different ways. In other words, discursive practices and material phenomena are intricately connected and mutually constituted, rather than being distinct or separable entities (Barad, 2003).

The agency axiom: Agency is not limited to humans, but rather is a potentiality that emerges through the interactions between all kinds of actors, both humans and nonhumans. The question of whether technology possesses agency has been widely debated within information systems studies, presenting at least three discernible perspectives. The first viewpoint, referred to as the humanistic perspective, contends that agency, defined as the 'capacity for action,' is an exclusive attribute of humans (Emirbayer and Mische, 1998; Giddens, 1984). Advocates of this stance posit

that the essence of agency resides in the unique human qualities of motivation, reflection, and rationalization. Thus, they posit that technology assumes a passive role, serving merely as a backdrop against which human actions and objectives unfold (Azad and King, 2008; Boudreau and Robey, 2005; Vaast and Walsham, 2005).

Conversely, the second perspective, labelled as 'post-humanistic', challenges the humanistic viewpoint for overlooking technology's potential agency and advocates for the equating of human and non-human agencies (Boudreau and Robey, 2005; Walsham, 2005). Drawing inspiration from the work of Latour (2007) and Callon (1990), this perspective asserts that human and non-human (e.g., technology) agencies are analogous, at least in semiotic terms. From this perspective, agency is not a fixed or inherent quality possessed by individual humans. Instead, it is a capacity that emerges from relational and shifting associations between actors, be they human or nonhuman (Latour, 2007). In other words, agency is not solely the result of an individual's internal qualities or characteristics, but it rather emerges in the context of social relationships and networks and is constantly shifting and evolving as those relationships change. In that sense, technologies possess agency, owing to their capacity to influence and exert an impact on the world (Orlikowski, 2005). Furthermore, it posits that technology exercises agency through its 'performativity,' (Barad, 2003; Pickering, 1995) actively shaping users' interactions with and perceptions of the world, often transcending users' direct control.

A third perspective, termed the 'soft post-humanistic,' acknowledges the presence of material agency while upholding its distinct and complementary status to human agency - the 'ability to form and realize goals' (Leonardi, 2011). Material agency, according to this, refers to the ability of non-human entities, such as technology, to act autonomously – without direct human intervention (Leonardi, 2011). For example, an email client can autonomously transmit a text message through an SMTP server, locate

the receiver's server, and deliver the message to the inbox, all while checking for spam and viruses, devoid of human involvement. While this perspective acknowledges technology's agency, it also addresses critiques of the post-humanistic approach, underscoring the differentiation between human and material agencies (Chae and Poole, 2005; Leonardi and Barley, 2008). Taylor (2001) proposes that these agencies influence disproportionately human actions, as humans have a 'head status,' while material technology assumes a "complement status." Significantly, from this perspective, human and material agencies are distinct and interdependently intertwined within organisations, generating technologies and routines through a process of imbrication – where overlapping patterns facilitate interdependent functioning (Ciborra, 2006; Sassen, 2006; Taylor, 2001). Neither human nor material agencies are independently pivotal, but their imbrication in specific sequences becomes transformative, resulting in the production, maintenance, or alteration of routines and technologies.

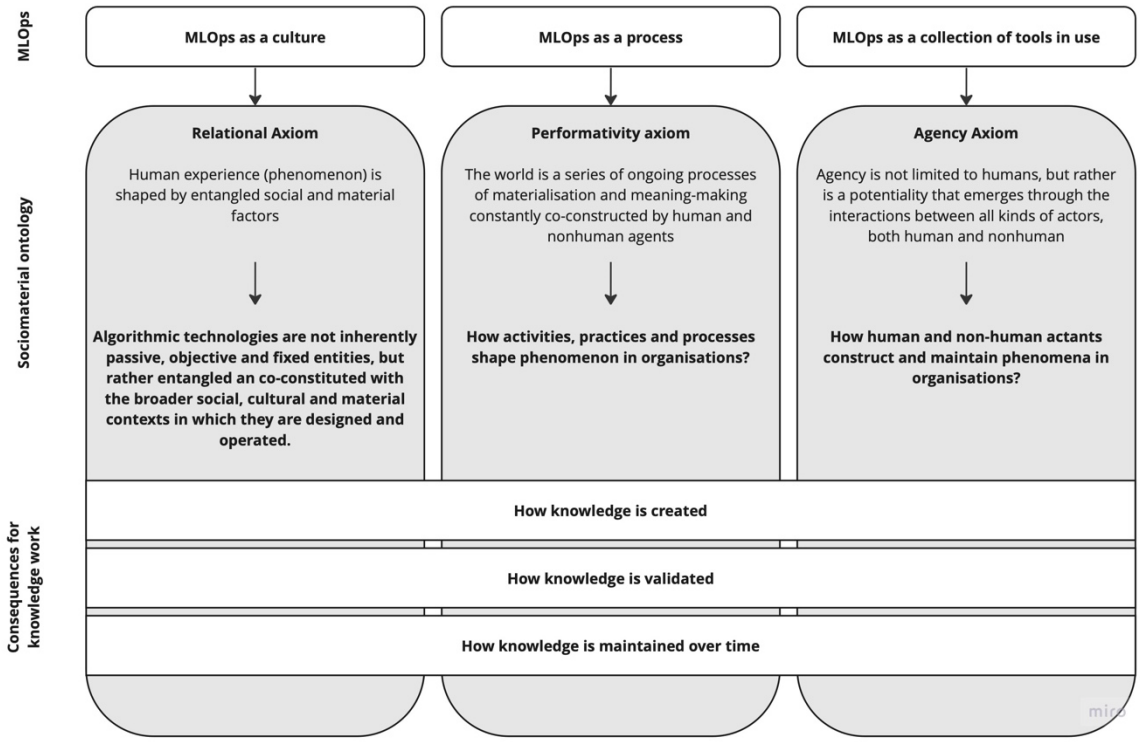
To explain why human and material agencies become imbricated, Leonardi (2011) draws on James Gibson's affordance theory (1986). This theory posits that the physical properties of an artefact exist independently of the user but acquire meaning contingent on the user's posture and behaviour. It emphasizes perception, highlighting that individuals do not simply apprehend an artefact's physical attributes (e.g., shape, colour, elasticity) but rather its affordances – the potential actions it facilitates or constrains. The perspective further emphasizes relationality, asserting that affordances manifest solely through interactions between individuals and the materiality of artefacts (Fayard and Weeks, 2007; Zammuto et al., 2007). Hence, individuals from distinct contexts, while interacting with the same artefact, may perceive varying affordances. This process of interpreting a technology through someone's own goals leads to perceiving different distinct possibilities for action (affordances and constraints), and different outcomes

(Leonardi, 2011, p.153). In this study, we adopt a soft post-humanistic perspective, acknowledging that technologies can act autonomously and thus, they have material agency.

5.3.2 The sociomateriality of MLOps

In this section, we situate the MLOps themes identified above in the context of organisations by drawing on a sociomaterial ontology. To that end, we draw on the axioms of relationality, performativity and agency described above. The first axiom of a sociomaterial ontology implies that machine learning algorithms are not inherently passive, objective and fixed entities, but rather entangled and co-constituted with the broader social, cultural, and material contexts in which they are designed and operated. Our theoretical apparatus is presented in Figure 5-1.

Figure 5-1 Theoretical apparatus for connecting common themes of MLOps with the effects of machine learning algorithms in knowledge work through a sociomaterial lens.



The first theme in MLOps we examined were the principles that shape the development culture enacted by data scientists embedded into organisational contexts, and thus we connected it to the relational axiom of sociomateriality. Understanding culture as relational means seeing it as a phenomenon that emerges, evolves, and is transmitted through social interactions among individuals and groups within a society. It is not a static construct, but rather a dynamic, ever-changing entity that is shaped by the actions and experiences of people in social contexts. Culture is relational because it exists in the relationships between individuals and groups, for example, between managers, data scientists, and users, and it is continuously being negotiated and contested through these relationships. Therefore, the culture promoted by the MLOps principles cannot be studied or understood in isolation from its social context but rather must be analysed concerning the social dynamics that give rise to it. Moreover, culture is relational because it is not a homogeneous entity, but rather a collection of shared meanings, beliefs, and practices that are specific to particular groups or communities.

These cultural expressions are created and maintained through social interactions within those groups and are shaped by the historical, social, and economic conditions that surround them. As such, it is essential to study MLOps principles concerning its social context and to recognize the diversity and complexity of cultural expressions within and between groups.

The second theme in MLOps we examined referred to the activities carried out in the development and operation of machine learning algorithms. We connect this theme to sociomateriality and argue that these activities are performative because they are not merely a sequence of steps, but rather they actively create or bring about the very thing they describe or enact. In other words, MLOps activities themselves have transformative effects on the actors or objects involved, shaping or producing them in ways that would not have been possible without the process. Performative processes are particularly salient in social and cultural contexts, where they are used to create and enact shared meanings, values, and practices. For example, a wedding ceremony is not simply a series of actions (such as exchanging rings or signing a certificate) that indicate the couple's legal or social status, but rather a performative process that creates and enacts the couple's new identity as a married couple. In this way, performative processes are not just descriptive, but also prescriptive. They do not just represent or reflect reality, but also actively shape it by bringing into being new social relations, identities, or practices. This is why performative processes are often seen as powerful tools for social change or transformation, as they can challenge or subvert existing norms, values, and power structures. In summary, a process is performative because it actively creates or brings about the very thing it describes or enacts and has a transformative effect on the actors or objects involved. Performative processes are particularly relevant in social and cultural

contexts, where they are used to create and enact shared meanings, values, and practices, and can be powerful tools for social change and transformation.

Finally, the third theme in MLOps we consider was the tools and artefacts deployed in the creation of machine learning. We link this to sociomateriality and consider that these artefacts and tools have agency. While we typically associate agency with human beings, it is also possible for non-human entities to possess a degree of agency. MLOps tools can have agency in several ways. First, they can enable or constrain human actions, shaping how we interact with the world. For example, data repositories allow data scientists to create machine learning models, but they also limit the data scientists' actions to the specific domain contained in the data. Second, MLOps artefacts can act autonomously, without human intervention, in certain contexts. For instance, training pipelines can be triggered automatically to retrain algorithmic models, changing the quality and availability of their predictions. Third, MLOps artefacts can embody or enact certain values, norms, or ideologies that shape how we think and act, in the same way, that smartphones, for example, shape norms of flexible work but also continual connectivity (Mazmanian et al., 2013). In all these cases, tools are not neutral or passive objects but rather possess a degree of agency that can shape human actions and outcomes. This does not mean that tools have human-like consciousness or intentionality, but rather that they are designed and used in ways that give them a certain degree of power and influence in the world. In other words, material means, such as code, data and models, are not merely tools that developers use to create new algorithmic technologies. Instead, they constitute the developer's identity and their actions.

In summary, enacting MLOps in organisations is a complex and dynamic process that must be understood with its social and cultural context, and the tools and artefacts involved have a significant impact on the development of machine learning algorithms.

We considered MLOps three interrelated themes, connecting them to the axioms of a sociomateriality ontology. The first theme concerns the culture that shapes the development of these algorithms, which is relational and specific to particular groups or communities. Drawing on the relational axiom of sociomateriality, the culture of MLOps is viewed as a dynamic and ever-changing entity that is shaped by social interactions among individuals and groups. This culture cannot be studied or understood in isolation from its social context but rather must be analysed in relation to the social dynamics that give rise to it. The second theme pertains to the performative nature of the activities involved in creating machine learning algorithms. These activities are not simply a sequence of steps but actively create or bring about the thing being described or enacted. As a result, performative processes are seen as powerful tools for social change or transformation, as they can challenge or subvert existing norms, values, and power structures. The third theme focuses on the tools and artefacts used in MLOps, which have agency and can shape human actions and outcomes. This agency is manifested in various ways, including enabling or constraining human actions, acting autonomously, and embodying or enacting certain values, norms, or ideologies. We now use this understanding to explore the effects of machine learning in knowledge work, considering both common epistemic and normative concerns identified in the literature.

5.4 Operational effectiveness implications of machine learning in knowledge work

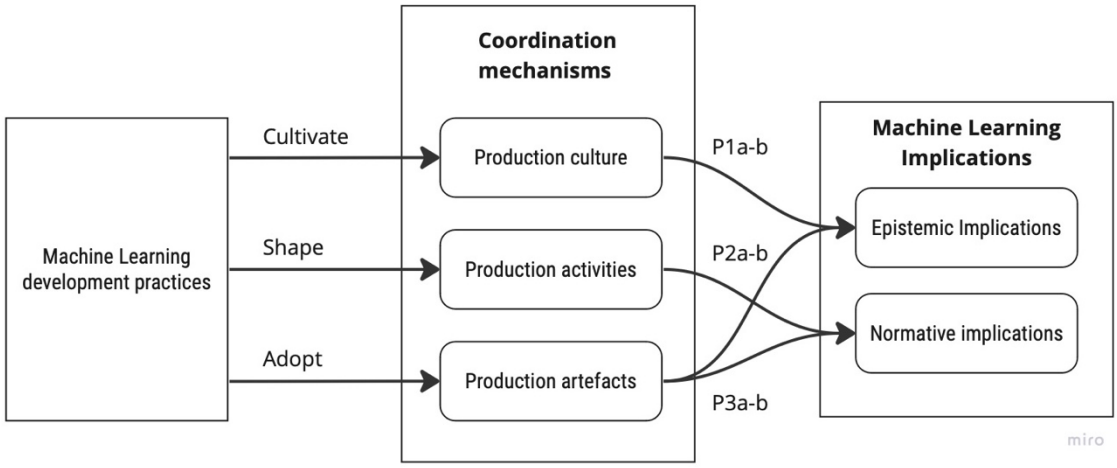


Figure 5-2 Conceptual model connecting a machine learning design and operations practice with its effectiveness implications.

The effects of MLOps practice on knowledge work are potentially far-reaching, particularly in terms of the creation, validation, and maintenance of the knowledge embedded in algorithmic-based processes, products and services. To make this concrete, we focus on how MLOps affects knowledge work, which requires primarily intellectual tasks. Knowledge work is one of the main contexts affected by machine learning algorithms, with applications now augmenting various fields such as medical diagnosis (Lebovitz et al., 2022), marketing (Chintalapati and Pandey, 2022) and legal advice (Sako et al., 2022). Although the eventual effects of machine learning in knowledge work may include changes in employment (Frey and Osborne, 2017), organisational structures and business models (Armour and Sako, 2020), these eventual effects depend on, first and foremost, the quality of the knowledge produced by algorithms. The foundation for creating value in knowledge work is, after all, knowledge itself. Thus, development and operations practices for creating and maintaining machine learning algorithms in general,

and MLOps in particular, have profound effects on the effects of these technologies in knowledge work.

In this section, we link each of the themes of MLOps, and the respective axiom, to the concerns of using algorithmic technologies in knowledge production, while also distinguishing between epistemic (Lebovitz et al., 2021) and normative concerns (Kane et al., 2021; Mittelstadt et al., 2016; Tsamados et al., 2021). Epistemic concerns relate to the quality and validation of the knowledge produced, such as whether algorithmic technologies generate evidence that blurs the line between fact and fiction. In contrast, normative concerns relate to questions of "what kind of world are we creating from embracing (and thereby tacitly ceding control to) machine learning" and other algorithmic technologies (Leavitt et al., 2021, p.2 original with machine learning abbreviated as ML).

Noteworthy is that we depart from the conventional usage of normative implications in Operations Management, which typically pertains to considerations of corporate policy, encompassing the guiding organisational mission, values, and norms. Instead, we use the term 'normative' in an ethical sense to refer to considerations of what is morally right or wrong. Our focus is on the consequences arising from the actions or judgments of algorithms. These include the generation of unfair outcomes through unjust discrimination (Barocas and Selbst, 2016; Green and Chen, 2019), as well as the algorithm's transformative effects on society (Mittelstadt et al., 2016; Tsamados et al., 2021). Importantly, normative implications are 'observer-dependent', as they depend on the individual's principles and values. Furthermore, these implications are independent of the quality of knowledge generated by algorithms, as they can arise even when the produced knowledge is based on conclusive, transparent, and well-founded evidence (Tsamados et al., 2021)

In the remainder of this section, we employ the three themes of MLOps, each with their respective sociomateriality axiom, to explain how development and operations practices shape the effects of algorithmic technologies in organisations. In doing so, we draw on a set of bodies of relevant literature that allow us to provide a rich understanding of the complex interplay between these themes and their impact on organisations. We also propose a set of propositions that encapsulate the key insights emerging from this review.

5.4.1 How the culture of MLOps shapes knowledge production

The writing on MLOps conceives knowledge as constructed by data scientists who enact a development culture rooted in shared epistemic beliefs as well as in principles of agility and continuous improvement. Two linked bodies of literature that provide rich insights into how relational cultures affect the adoption of new technology-enabled practices are those on 'epistemic cultures' (Knorr Cetina, 1999) and 'epistemic stances' (Fayard et al., 2016). Epistemic cultures refer to 'arrangements and mechanisms - bonded through affinity, necessity, and historical coincidence – which, in a given field, make up how we know what we know'. In other words, they are cultures that "create and warrant knowledge" (Knorr Cetina, 1999, p.1). In contrast to epistemic cultures, the notion of epistemic stances highlights the micro-level dynamics and enactments toward pursuing knowledge (Chakravartty, 2011). For instance, data scientists enacting agile product development share the same epistemic culture as data scientists enacting MLOps development, which is a product of their training, tools, and socialization. Yet they differ in their epistemological stance, as the organisation, its principles and practices differ. The constructs of epistemic cultures and stances are premised on the idea that the generation of knowledge hinges on particular attitudes and beliefs about the nature of

the world, as well as the modes of inquiry employed to investigate the world, thereby positing culture as a central, driving force in knowledge work.

To translate these ideas into the context of the culture of MLOps, we draw on the relational axiom of sociomateriality outlined earlier. It leverages the idea that the interactions between actors, and their nature, construct phenomena in organisations. Here the emphasis is on the relationships that may reshape the meaning and interpretation of their production. In the context of the culture of MLOps, this might involve understanding how relationships between the data scientists, managers, and users in the organisation may affect their epistemic stance and thus, the generation and warrant of knowledge.

The epistemic culture of data scientists is shaped by their multi-disciplinary backgrounds and characterized by an emphasis on empirical evidence, an inclination to adopt new methods, a proclivity to collaboration, and attention to detail. Data scientists draw on various disciplines, including computer science, statistics, mathematics, and domain-specific knowledge to work on and transform large datasets into machine learning models and predictions. To do so, data scientists rely on historical evidence and statistical analysis to support their findings and make decisions; they constantly seek out new methods and technologies to improve their analysis and insights; they often work across interdisciplinary teams; and they are meticulous in their work, paying close attention to detail to ensure the accuracy of their analysis. Overall, the epistemic culture of data scientists is characterized by a strong emphasis on rigorous empirical analysis, collaboration, and continuous learning, all of which are necessary to extract insights from large and complex datasets.

However, when data scientists are embedded into organisations, their epistemic stance gets modified by other actors, including managers and users. A culture of MLOps

highlights the values of agility and continuous improvement. Nevertheless, not all deployments of MLOps show the same *fidelity* – to use Ansari et al.'s (2010) language. Some implementations remain closer to one attribute rather than the other. In many cases, organisations prioritise agility and promote cultures of production. The focus on agility encourages a focus on delivering products as fast as possible, creating deadline pressures that legitimise errors and mistakes (Sculley et al., 2014), resulting in a change in data scientists' epistemic stance. The attributes of this new culture affect data scientists' attitudes towards the production of knowledge. In this new stance, data scientists pay less attention to detail and focus less on practices for evaluating knowledge. Implementations of an MLOps culture that focuses on rapid production, thus leading to promoting an epistemic stance that normalises deviance. Vaughan (1996) offers an example in the field of engineering where the normalisation of deviance ultimately became the catalyst for the Challenger space shuttle disaster. Extrapolating to data science, it is evident how the promotion of agile values as lived in the organisation can recreate similar situations in which the data scientists make mistakes that lead to algorithmic technologies producing catastrophic consequences for organisations.

When organisations adopt the MLOps based on the principles of continuous improvement – as opposed to sole agility - they legitimise the permissions to correct mistakes. This creates an environment where data scientists are more likely to adopt an epistemic stance that aligns with their epistemic culture, promoting a thorough evaluation of knowledge and attention to detail. Data scientists are more likely to enact an epistemic stance that is closer to their epistemic culture, and thus pay more attention to detail, and the evaluation of knowledge. These implantations will encourage data scientists to pay attention to detail and to evaluate algorithmic technology outputs, assessing whether the output quality is appropriate.

The effects of the MLOps culture on the consequences of algorithmic technologies for organisations are potentially profound. A relational perspective emphasizes the active and influential role of the relationships, and their nature in shaping data scientists, and thus the algorithmic technologies produced. For an organisation developing algorithmic technologies, the culture has a significant effect on affecting the epistemic consequences of algorithmic technologies. However, it is unlikely to affect the normative consequences of algorithmic technologies While the development culture modifies the epistemic stance of developers, it does not affect their onto-epistemological assumptions. In other words, data scientists, still have an empirical predisposition to understand the world and truth as the result of historical patterns, resembling Alfred North Whitehead "dustbowl empiricism". As Suddaby (2014) warns us, such a view can be dangerous as it "… discourage[s] researchers from asking fundamental questions about the assumptions that underpin knowledge and the methods used to acquire knowledge" (p.408). Consider, for instance, the case of developing a machine learning model to recommend which candidates to hire based on historical data. Data scientists are likely to perceive historical data as an objective and truthful representation of which candidates are more likely to be successful in the organisation. Because of this onto-epistemological view, data scientists are less likely to consider that this data might have been biased and that hiring may be also a product of the recruiters' biases. Thus, even if they had the time to explore the investigate the data, this pattern of failure would not have been identified. Thus, we argue that changing the development value of the organisations is unlikely to change the organisations' ability to deal with normative concerns caused by algorithmic technologies – or well, at least in the immediate term. Thus, we propose:

Proposition 1a: A machine learning design and operation practice that cultivates a production culture based on continuous improvement, rather than agility, reduces the *epistemic* errors of the resultant machine learning outputs.

Proposition 1b: A machine learning design and operation practice that cultivates a production culture based on continuous improvement, rather than agility, has no effect on the *normative* errors of the resultant machine learning outputs.

To effectively test these propositions, a comparative analysis can be undertaken between two distinct groups: perpetrator organisations, those linked to incidents associated with machine learning, and non-perpetrators, comprising machine learning development teams without any such incidents. This comparative examination can be conducted through a comprehensive survey administered to machine learning development teams engaged in various projects. The survey should encompass targeted questions that centre on the prevalence of agility and continuous improvement within their respective organisational cultures. Furthermore, it should collect detailed information about any machine learning incidents encountered by the involved companies.

To facilitate the analysis, incidents can be systematically categorized into two types: normative and epistemic, employing clearly defined definitions for each category. To ensure impartial and rigorous coding of incidents, the involvement of multiple coders is highly recommended. Through rigorous analysis, researchers can explore potential associations between organisational cultures that emphasize continuous improvement or agility and the occurrence of incidents in machine learning projects. This could be done,

for example, by calculating the tetrachoric correlation between the type of culture and the type of machine learning incident.

In testing these propositions, researchers could shed crucial light on the pivotal role of organisational culture in fostering responsible machine learning practices. Additionally, they could identify potential risk factors that contribute to incidents in this domain. By examining the contrasting experiences of perpetrator and non-perpetrator organisations, further research on this stream holds promise in advancing our understanding of effectiveness considerations in machine learning and promoting the adoption of responsible practices in the field.

5.4.2 How practices in MLOps shape knowledge outputs

We now connect the second key theme in writing about MLOps, activities and processes, with the consequences of algorithmic technologies for knowledge work. An orientation towards activities and processes emphasises the value of day-to-day micro-actions in constructing and maintaining algorithmic technologies and their effects on organisational life. As discussed above, this theme has a lot in common with writing on sociomateriality that draws attention to the performativity of actions and discourses in shaping the world. We thus, focus on the performativity of the activities and discourses involved in the adoption of MLOps to explore how they shape epistemic and normative consequences of algorithmic technologies for organisations. To explore this in further detail, we consider the notion of adaptation of MLOps in organisations (Ansari et al., 2010), and how it affects the consequences of machine learning in organisations.

First, we focus on full adaptation, whereby organisations appropriate the activities and processes of MLOps and its core principles The processes and activities promoted in MLOps reveal a strong emphasis towards addressing epistemic concerns, but not towards addressing normative concerns. For example, some MLOps activities such as cleaning

and preparing the data, testing the machine learning model and monitoring its performance over time, aim to ensure the knowledge produced by algorithmic technologies complies with customer's specifications or imposed conditions. These activities, thus, minimise the epistemic concerns of algorithmic technologies by verifying the quality of the knowledge produced. Nevertheless, they are less concerned with knowledge validation, the assurance that the knowledge generated meets customers' and stakeholders' needs. And so, they fail to address the normative consequences of algorithmic technologies. The practices do not prompt developers to contemplate if the data leads to argue that when correctly applied the practices of MLOps can indeed transform the consequences of algorithmic technologies in organisational life. Thus, we propose:

Proposition 2a: A machine learning design and operation practice that establishes formal production activities for validation and testing, and reduces the normative errors of the resultant machine learning outputs, when the production culture is anchored in continuous improvement.

We now direct our attention to low-fidelity adaptations of MLOps in which organisations appropriate activities and processes of MLOps without embracing its core principles. Such adoptions emerge from rhetorical -and not substantive – implementations of MLOps. We argue that low-fidelity adaptations that involve copying the processes and activities of MLOps without changing the development culture are unlikely to assist organisations in addressing the epistemic or normative concerns of algorithmic technologies. To explain this, consider the disparity between rhetoric and the

actual performativity of practices noted in the diffusion of other technical practices (Zbaracki, 1998). When organisations adopt MLOps, they emulate the processes and activities of MLOps as a package of scripts. However, the actual execution of these activities may diverge due to managerial ignorance, technical intimidation (Zbaracki, 1998), and production pressures. Managers with little knowledge of MLOps may hold unrealistic expectations of its tools, principles, and practices. Coupled with the technical intricacies of MLOps, which involve statistics and machine learning principles, managers may cling to a magical view of the technology in the guise of artificial intelligence, failing to grasp the realities of MLOps. Furthermore, production pressures may compel data scientists to bypass some of the development activities, diminishing its effectiveness in addressing epistemic concerns. This results in an adoption form that is legitimized, but wherein developers may omit some of the actions established in the sequence or implement them incompletely.

Proposition 2b: A machine learning design and operation practice that establishes formal production activities for validation and testing, increases the epistemic errors of the resultant machine learning outputs when the production culture is not anchored in continuous improvement.

To test the second set of propositions, we recommend adopting a similar approach as used for testing the first set. This involves comparing perpetrator organisations and non-perpetrator organisations by surveying teams that develop machine learning-based solutions. The survey should cover questions related to two key aspects: (a) Prevalence of agility and continuous improvement within their respective organisational cultures; and (b) The establishment of formal activities for testing and validating the outputs of machine learning. Researchers can then employ a tetrachoric correlation between the

existence of formal testing mechanisms and the type of incident associated with the machine learning system. This analysis should also control for the type of culture present within each organisation.

5.4.3 How the collection of tools and artefacts of MLOps shapes knowledge outputs

The writing on MLOps also discusses the tools and artefacts that developers use to construct algorithmic-based products and services. Through a sociomaterial perspective, these tools and artefacts are understood not as passive, neutral objects, but rather as agentic entities that can shape and impact the normative and epistemic consequences of algorithmic technologies. Among these tools and artefacts are data repositories, algorithmic models, pipelines, and automated training and testing scripts. Such tools and artefacts can have a significant impact on the functioning and outcomes of algorithmic technologies. For instance, data repositories and labels may contain biases and inaccuracies that can be perpetuated in algorithmic models, particularly in cases involving subjective assessments. The adages "bias in, bias out", and "garbage in, garbage out" come to mind.

Two poignant illustrations of how data repositories and labels can give rise to normative and epistemic concerns are Amazon's discriminatory hiring algorithm (Dastin, 2018) and Lebovitz et al's (2021) investigation of inaccurate algorithms in healthcare. The Amazon case underscores the role of data artefacts in shaping normative concerns. In developing an algorithm to screen candidates, Amazon relied on historical data spanning a decade from applicants' submissions. However, since the historical data primarily comprised male applicants, the algorithm learnt patterns that advantaged male candidates while discriminating against female candidates. The historical data was not

gender-neutral, and as such it shaped the hiring algorithm and its recommendations, resulting in normative concerns that ultimately led Amazon to scrap the tool. Similarly, Lebovitz et al. (2021) shows an example in healthcare of how data and labels contributed to epistemic concerns. Specifically, a radiology department developed algorithms to analyse medical images and classify them into illnesses. Experts assigned labels (e.g., brain tumour) to each image to enable the algorithm to learn how to segment images into illnesses. However, the study reveals that these labels were inherently incomplete - failed to capture the actual physician's practice - subject to multiple interpretations and, ultimately, somewhat unreliable. Thus, even though professionals had created the labels, the resulting algorithms were plagued by epistemic concerns, increasing the uncertainty behind the medical diagnosis.

Depending on the type of technology, algorithmic models may further exacerbate epistemic concerns as their outputs are inscrutable. The use of deep learning architectures is a prime example of this issue. These architectures are associated with exceptionally high levels of performance, but they do not explain how the algorithm arrived at the solution. Unsupervised learning algorithms may also present similar challenges, collapsing a set of data points into an indistinguishable form. On the other hand, supervised learning techniques that do not use deep learning architectures, such as regression and classification algorithms, are more scrutable. Here, one could analyse the predictions these models produce by examining the coefficients of the regression or by tracing the decision points in three based models. Because data and model artefacts are common in early-stage implementations of MLOps (Symeonidis et al., 2022), we suggest:

Proposition 3a: A machine learning design and operation practice that adopts automation artefacts increases the prevalence of epistemic and normative concerns in the short term.

Not all artefacts in MLOps exacerbate epistemic and normative concerns of algorithmic technologies. Records of model versions, as well as automation artefacts such as pipeline, and training and testing scripts can play a pivotal role in redressing biases and improving the veracity of the models. Model version records, for example, enable developers to easily revert to a previous iteration of the algorithm in the event of inaccuracies, errors, or biases. This is a crucial step in mitigating epistemic and normative concerns, as it allows developers to quickly identify and address potential issues. Additionally, pipelines, automatic training, and testing scripts can further reduce the likelihood of such concerns arising in MLOps. These automation artefacts not only accelerate the production of algorithmic technologies but also enable developers to quickly make necessary corrections. Automatic testing scripts facilitate swift detection of potential issues, while automatic training scripts enable continual optimization of the model's performance through retraining on new data. As a result, organisations that adopt such approaches can enhance the efficiency, reliability, and effectiveness of their MLOps processes. Because automation artefacts are usually associated with organisations with a more mature MLOps capability (Symeonidis et al., 2022), we propose:

Proposition 3b: A machine learning design and operation practice that adopts automation artefacts decreases the prevalence of epistemic and normative concerns in the long term.

To test the third set of propositions, we recommend a longitudinal survey in organisations actively developing machine learning solutions. This survey should focus on two main aspects: (a) Tracking the prevalence of issues associated with their machine learning solutions at different points in time; (b) assessing the level of automation involved in developing these solutions at specific time intervals. By conducting this longitudinal survey, we can gain valuable insights into how effectiveness issues evolve over time in the context of machine learning development and the impact of automation on these developments.

5.5 Conclusion

In this article, we explore the intricate interplay between development and operations practices in algorithmic technologies and their operational effectiveness implications. Our investigation reveals the need for nuanced and contextually grounded research in this realm, challenging the notion of algorithms as inherently neutral. Instead, we highlight their susceptibility to social and political influences in their deployment contexts, aligning with recent debates (Anthony et al., 2023; Bailey and Barley, 2020; Bailey et al., 2019; Faraj and Pachidi, 2021).

Specifically, this paper makes a significant contribution by providing a conceptual model that connects the design and operation of machine learning with its operational effectiveness implications. The analysis centres on MLOps, considering it as a representative instance of machine learning practices in design and operation. By distilling key characteristics of such practices—production culture, production activities,

and production artefacts—we offer a model that sheds light on how operational practices influence coordination mechanisms and subsequently impact the operational effectiveness implications of machine learning.

Although our model can be generalized beyond MLOps, it is crucial to underscore the fundamental distinctions between MLOps and other practices, like DevOps. In DevOps, data scientists are responsible for creating the product, entailing tasks such as writing code for deployment and designing evaluation tests. However, MLOps introduces a paradigm shift, where machine learning generates the code itself to accomplish tasks. This shift renders the underlying rules and mechanisms opaque, particularly in complex deep learning architectures. Consequently, comprehending the ML program and devising appropriate test cases to ensure accurate results from the code becomes a challenging task for data scientists.

At first glance, it may seem that the main difference between MLOps and DevOps lies solely in the production artefacts used (as MLOps relies on machine learning while DevOps does not). However, assuming this would be incorrect. The dissimilarities extend beyond artefacts, as the production culture, artefacts, and activities of a practice are interconnected and mutually influence one another. This paper argues that artefacts possess material agency, enabling them to act autonomously without human intervention. As managers and data scientists endeavour to align their production goals with the material agency of machine learning, they construct perceptual affordances and constraints that, in turn, affect cultural norms and production activities.

To put it simply, when managers perceive machine learning to accelerate product creation, they are likely to reinforce norms of agility that shape the production culture, possibly leading to the avoidance of thorough testing and support activities. This suggests that design and operation practices utilizing different artefacts are likely to cultivate

diverse cultures and promote varying activities. Understanding these interconnected aspects is essential for navigating machine learning effectively.

By recognizing the complexity of the interplay between algorithms and organisational practices, we can design and implement algorithms more thoughtfully, aligning them with organisational goals and mitigating unintended negative consequences. This direction of research offers the potential to advance both theoretical understanding and practical insights for organisations dealing with algorithmic technologies. These implications extend to researchers, practitioners, and policymakers, aiming to promote the responsible use of algorithms in organisations and contribute to societal well-being (Berente et al., 2021).

5.5.1 Theoretical contributions: From neutrality to complexity

The first implication of our study is to highlight the need for research on the interplay between development and operations practices in shaping algorithmic technologies. Historically, management and economic research have prioritized studies examining how technology impacts organisations and their practices (e.g. Brynjolfsson et al., 2021). Indeed, we concur that algorithmic technologies play an essential role in transforming organisational life and labour by, for example, promoting the emergence of new roles, such as prompt engineers and translators of algorithmic predictions (Waardenburg et al., 2021). Yet, these studies have neglected the equally important notion that algorithmic technologies are shaped by organisations themselves. To address this issue, our study theorizes how established organisational practices shape algorithmic technologies, the knowledge embedded in products and services, and their effects on organisations. We propose that changing the principles of developing and operating algorithmic technologies has the potential to improve epistemic concerns. We suggest three specific directions for future research in this stream: Researchers interested in the role of development and

operations practices should focus on the relational development culture, including epistemic beliefs, the rhetoric and reality of organisational principles, and the enacted epistemic stances involved in the design of algorithmic technologies, and that arise from interactions between developers, managers, and users. Moreover, our relational perspective emphasizes the need for studying how the tensions in the development culture emerge and are resolved across different types of knowledge work organisations.

Our study also suggests that activity-oriented studies can make an important contribution to our understanding of the relationship between development and operation practices and the effects of algorithmic technologies in organisations. We propose that relying solely on principles is inadequate in addressing the normative concerns associated with algorithmic technologies and that organisations need to adapt their practices to effectively mitigate these concerns, for example, by incorporating specific activities to audit algorithmic technologies (Floridi et al., 2022). Future research should focus on identifying failures resulting from inadequate enactment of development and operation activities, and exploring how managers and developers respond to them through strategic actions. Furthermore, researchers could investigate the extent to which MLOps substitute or complement efforts to create explainable artificial intelligence (e.g. Hafermalz and Huysman, 2022). It remains unclear whether epistemic and normative concerns are better addressed through practices that scrutinize predictions, akin to the approaches of knowledge workers (Lebovitz et al., 2021; Lebovitz et al., 2022), or through practices that emphasize rapid feedback and corrections, as in MLOps. Similarly, future research should examine specific activities that modify knowing practices (Anthony, 2018) and the regime of knowing (Pachidi et al., 2021) to ensure effective and productive adoption of algorithms, justifying continued support.

A third research implication concerns the materiality of algorithmic technologies. Consistent with research on sociomateriality (Cecez-Kecmanovic et al., 2014; Jones, 2014; Orlikowski and Scott, 2008, 2015a) and boundaries (Barrett et al., 2012; Langley et al., 2019) we have argued that matter matters. Our main point is that artefacts may have an important role in shaping the consequences of algorithmic technologies in organisational life. The link we establish between sociomateriality and MLOps, suggests that artefacts are not neutral but complex objects that shape and are shaped by organisations, creating feedback loops with profound normative and epistemic effects. Future research should thus focus on understanding how organisations deal with the paradox that the artefacts of MLOps pose for the development of algorithmic technologies – namely, that these artefacts are more likely to produce models that violate social and ethical norms in the short term, but still enable organisations to produce robust models in the long term.

5.5.2 Managerial implications

Our conceptual model has important practical implications for organisations attempting to develop effective algorithmic technologies: First and foremost, organisations need to change the principles underpinning the design and operations of machine learning by replacing principles of *agility* with those of *continuous improvement*. Practically speaking, this change should involve creating spaces for socializing those principles among the team members developing algorithmic-based products and services. This may involve conversations where managers and developers articulate their understanding of such principles while agreeing upfront on how to negotiate tensions between the pressure to launch new products and improving and maintaining existing ones.

Once organisations establish continuous improvement as a guiding principle for their algorithmic technologies, a shift in the production and maintenance activities is

required. Our analysis suggests that this can be facilitated by adopting workflows of steps used, for example, in MLOps. These procedures ought to be complemented with specific activities designed to identify and address biases, inequalities; and other normative issues in the context of deployment. Such activities include integrating specific steps to check for correlations between outputs and protected and sensitive variables, validating whether the data used for machine learning perpetuates, deliberately or invertedly, social bias, and encouraging knowledge workers to provide feedback on the algorithm recommendations. These mechanisms help to make salient epistemic and normative concerns caused by algorithmic technology recommendations, and as such create machine learning-based products and services that truly add value to organisations.

Enabling an effective monitoring and continuous evaluation of the knowledge generated by algorithmic technologies requires a range of skills, structures, and processes. To facilitate such evaluation, organisations require, at a minimum, practitioners with expertise in statistical inference and a solid understanding of the principles of machine learning algorithms. These practitioners might act as knowledge brokers, bridging the gap between users and developers, building trust, and promoting collaboration to transform issues, whether inaccuracies or biases, into actionable corrections. The significance of a continuous evaluation also highlights the need to integrate automation wherever possible, to enable speedy identification and correction of issues. If anchored in the right culture, automation can encourage organisational actors to voice concerns about algorithmic problems constructively, without assigning blame, thus enhancing the quality of the products and services generated. This approach would promote a culture of accountability for the quality of algorithms and encourage change across the organisation.

Finally, it is crucial to acknowledge that the generic approach outlined in this study may not be appropriate for all types of organisations. The purpose of this study has been to advance a theoretical understanding of the role of development and operations practices in shaping the impact of machine learning in organisations. By applying the

concepts presented here, organisations can increase the likelihood of rapidly producing algorithmic-based products and services, while also mitigating normative and epistemic concerns. Nonetheless, it is important to recognize that not all organisations have the same level of ambition: some organisations may symbolically adopt machine learning, to gain legitimacy from relevant stakeholders, or in an experimental manner, seek to evaluate the feasibility of new technologies. In such situations, the costs of acquiring orchestration platforms, and developing the required skills, culture, and collaboration structures may not be justified. Consequently, the significance of meticulously assessing the stakes and scale involved in implementing machine learning in organisational contexts should not be underestimated.

References

Ahmad, M.O., Dennehy, D., Conboy, K. & Oivo, M. (2018). Kanban in software engineering: A systematic mapping study. *Journal of Systems and Software, 137*, 96-113.

Ahmad, M.O., Markkula, J. & Oivo, M. (2013). *Kanban in software development: A systematic literature review*, 2013 39th Euromicro conference on software engineering and advanced applications. IEEE, pp. 9-16.

Alla, S. & Adari, S.K. (2021). 'What is MLOps?', *Beginning MLOps with MLFlow*. Apress, Berkeley, pp. 79-124.

Ansari, S.M., Fiss, P.C. & Zajac, E.J. (2010). Made to fit: How practices vary as they diffuse. *Academy of Management Review, 35* (1), 67-92.

Anthony, C. (2018). To question or accept? How status differences influence responses to new epistemic technologies in knowledge work. *Academy of Management Review, 43* (4), 661-679.

Anthony, C. (2021). When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative Science Quarterly, 66* (4), 1173-1212.

Anthony, C., Bechky, B.A. & Fayard, A.-L. (2023). "Collaborating" with AI: Taking a System View to Explore the Future of Work. *Organization Science, Forthcoming*

Armour, J. & Sako, M. (2020). AI-enabled business models in legal services: From traditional law firms to next-generation law companies? *Journal of Professions and Organization, 7* (1), 27-46.

Ashforth, B.E., Gioia, D.A., Robinson, S.L. & Treviño, L.K. (2008). Re-viewing organizational corruption. *Academy of Management Review, 33* (3), 670-684.

Azad, B. & King, N. (2008). Enacting computer workaround practices within a medication dispensing system. *European Journal of Information Systems, 17* (3), 264-278.

Bailey, D.E. & Barley, S.R. (2020). Beyond design and use: How scholars should study intelligent technologies. *Information and Organization, 30* (2), 100286.

Bailey, D.E., Faraj, S., Hinds, P., von Krogh, G. & Leonardi, P.M. (2019). Special issue of organization science: Emerging technologies and organizing. *Organization Science, 30* (3), 642-646.

Barad, K. (2003). Posthumanist performativity: Toward an understanding of how matter comes to matter. *Signs: Journal of women in culture and society, 28* (3), 801-831.

Barad, K. (2007). *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Duke University Press, Durham.

Barley, S.R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly, 31* (1), 78-108.

Barocas, S. & Selbst, A.D. (2016). Big data's disparate impact. *California law review, 104* (3), 671-732.

Barrett, M., Oborn, E., Orlikowski, W.J. & Yates, J. (2012). Reconfiguring boundary relations: Robotic innovations in pharmacy work. *Organization Science, 23* (5), 1448-1466.

Beane, M. & Orlikowski, W.J. (2015). What difference does a robot make? The material enactment of distributed coordination. *Organization Science, 26* (6), 1553-1573.

Beck, K. (1999). *Extreme programming explained: embrace change*, 2nd ed. Addison-Wesley Professional, Boston.

Beck, K. & Fowler, M. (2001). *Planning extreme programming.* Addison-Wesley Professional,

Benbya, H., Davenport, T.H. & Pachidi, S. (2020). Special issue editorial: Artificial intelligence in organizations: Current state and future opportunities. *MIS Quarterly Executive, 19* (4), IX-XXI.

Bengio, Y., Lamblin, P., Popovici, D. & Larochelle, H. (2006). Greedy layer-wise training of deep networks. *Advances in Neural Information Processing Systems, 19*

Bengio, Y. & LeCun, Y. (2007). Scaling learning algorithms towards AI. *Large-scale kernel machines, 34* (5), 1-41.

Berente, N., Gu, B., Recker, J. & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly, 45* (3), 1433-1450.

Bessant, J., Caffyn, S., Gilbert, J., Harding, R. & Webb, S. (1994). Rediscovering continuous improvement. *Technovation, 14* (1), 17-29.

Bhuiyan, N. & Baghel, A. (2005). An overview of continuous improvement: from the past to the present. *Management decision, 43* (5), 761-771.

Boehm, B.W. (1988). A spiral model of software development and enhancement. *Computer, 21* (5), 61-72.

Boudreau, M.-C. & Robey, D. (2005). Enacting integrated information technology: A human agency perspective. *Organization Science, 16* (1), 3-18.

Brynjolfsson, E., Wang, C. & Zhang, X. (2021). The economics of IT and digitization: Eight questions for research. *MIS Quarterly, 45* (1), 473-477.

Burgess, T.F. (1994). Making the leap to agility: Defining and achieving agile manufacturing through business process redesign and business network redesign. *International Journal of Operations & Production Management, 14* (11), 23-34.

Callon, M. (1990). Techno-economic networks and irreversibility. *The sociological review, 38* (S1), 132-161.

Callon, M. & Muniesa, F. (2005). Peripheral vision: Economic markets as calculative collective devices. *Organization Studies, 26* (8), 1229-1250.

Cao, L. (2022). AI in finance: challenges, techniques, and opportunities. *ACM Computing Surveys (CSUR), 55* (3), 1-38.

Cecez-Kecmanovic, D., Galliers, R.D., Henfridsson, O., Newell, S. & Vidgen, R. (2014). The sociomateriality of information systems. *MIS Quarterly, 38* (3), 809-830.

Chae, B. & Poole, M.S. (2005). The surface of emergence in systems development: agency, institutions, and large-scale information systems. *European Journal of Information Systems, 14* (1), 19-36.

Chakravartty, A. (2011). A puzzle about voluntarism about rational epistemic stances. *Synthese, 178* (1), 37-48.

Chau, M., Li, T.M., Wong, P.W., Xu, J.J., Yip, P.S. & Chen, H. (2020). Finding People with Emotional Distress in Online Social Media: A Design Combining Machine Learning and Rule-Based Classification. *MIS Quarterly, 44* (2), 933-955.

Chintalapati, S. & Pandey, S.K. (2022). Artificial intelligence in marketing: A systematic literature review. *International Journal of Market Research, 64* (1), 38-68.

Chuang, H.H.C., Chou, Y.C. & Oliva, R. (2021). Cross-item learning for volatile demand forecasting: An intervention with predictive analytics. *Journal of Operations Management, 67* (7), 828-852.

Ciborra, C. (2006). Imbrication of representations: Risk and digital technologies. *Journal of Management Studies, 43* (6), 1339-1356.

Collins, J.D., Uhlenbruck, K. & Rodriguez, P. (2009). Why firms engage in corruption: A top management perspective. *Journal of Business Ethics, 87* (1), 89-108.

Costas, J. & Fleming, P. (2009). Beyond dis-identification: A discursive approach to self-alienation in contemporary organizations. *Human relations, 62* (3), 353-378.

Curchod, C., Patriotta, G., Cohen, L. & Neysen, N. (2020). Working for an algorithm: Power asymmetries and agency in online work settings. *Administrative Science Quarterly, 65* (3), 644-676.

Dastin, J. (2018). *Amazon scraps secret AI recruiting tool that showed bias against women*. Available at: (Accessed: April 4th 2023)

DeSanctis, G. & Poole, M.S. (1994). Capturing the complexity in advanced technology use: Adaptive structuration theory. *Organization Science, 5* (2), 121-147.

Dove, R. (1991). The 21s Century Manufacturing Enterprise Strategy. *Iacocca Institute, Lehigh University, AD-A257176,*

Dyck, A., Penners, R. & Lichter, H. (2015). *Towards definitions for release engineering and DevOps*, 2015 IEEE/ACM 3rd International Workshop on Release Engineering. IEEE, pp. 3-3.

Emirbayer, M. (1997). Manifesto for a relational sociology. *American Journal of Sociology, 103* (2), 281-317.

Emirbayer, M. & Mische, A. (1998). What is agency? *American Journal of Sociology, 103* (4), 962-1023.

Faraj, S. & Pachidi, S. (2021). Beyond Uberization: The co-constitution of technology and organizing. *Organization Theory, 2* (1), 1-14.

Faraj, S., Pachidi, S. & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization, 28* (1), 62-70.

Fayard, A.-L., Gkeredakis, E. & Levina, N. (2016). Framing innovation opportunities while staying committed to an organizational epistemic stance. *Information Systems Research, 27* (2), 302-323.

Fayard, A.-L. & Weeks, J. (2007). Photocopiers and water-coolers: The affordances of informal interaction. *Organization Studies, 28* (5), 605-634.

Floridi, L., Holweg, M., Taddeo, M., Amaya Silva, J., Mökander, J. & Wen, Y. (2022). capAI-A Procedure for Conducting Conformity Assessment of AI Systems in Line with the EU Artificial Intelligence Act. *Available at SSRN 4064091,*

Fowler, M. & Highsmith, J. (2001). The agile manifesto. *Software development, 9* (8), 28-35.

Frey, C.B. & Osborne, M.A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114* (2017), 254-280.

Garg, S., Pundir, P., Rathee, G., Gupta, P., Garg, S. & Ahlawat, S. (2021). *On continuous integration/continuous delivery for automated deployment of machine learning models using mlops*, 2021 IEEE fourth international conference on artificial intelligence and knowledge engineering (AIKE). IEEE, pp. 25-28.

Gibson, J.J. (1986). *The ecological approach to visual perception*. Erlbaum, London.

Giddens, A. (1984). *The constitution of society : outline of the theory of structuration*. Polity, Cambridge.

Gift, N. & Deza, A. (2021). *Practical MLOps: Operationalizing machine learning models*. O'Reilly Media, Inc., Online.

Goldratt, E.M. (1997). *Critical chain*. North River Press, Great Barrington, Mass.

Goldratt, E.M. & Cox, J. (1984). *The goal: excellence in manufacturing*. North River Press, Croton-on-Hudson, N.Y.

Goodfellow, I., Bengio, Y. & Courville, A. (2016). *Deep learning*. The MIT Press, Cambridge, Massachusetts.

Google (2020). *MLOps: Continuous delivery and automation pipelines in machine learning*. Available at: https://cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning (Accessed: August 17th 2022)

Granlund, T., Kopponen, A., Stirbu, V., Myllyaho, L. & Mikkonen, T. (2021). *Mlops challenges in multi-organization setup: Experiences from two real-world cases*, 2021 IEEE/ACM 1st Workshop on AI Engineering-Software Engineering for AI (WAIN). IEEE, pp. 82-88.

Green, B. & Chen, Y. (2019). *Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments*, Proceedings of the conference on fairness, accountability, and transparency, pp. 90-99.

Greve, H.R., Palmer, D. & Pozner, J.E. (2010). Organizations gone wild: The causes, processes, and consequences of organizational misconduct. *Academy of Management Annals, 4* (1), 53-107.

Hafermalz, E. & Huysman, M. (2022). Please Explain: Key Questions for Explainable AI research from an Organizational perspective. *Morals & Machines, 1* (2), 10-23.

Hewage, N. & Meedeniya, D. (2022). Machine learning operations: A survey on MLOps tool support. *arXiv preprint arXiv:2202.10169,*

Hinton, G.E., Osindero, S. & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural computation, 18* (7), 1527-1554.

Holweg, M., Davies, J., De Meyer, A., Lawson, B. & Schmenner, R.W. (2018). *Process theory: The principles of operations management*. Oxford University Press, Oxford.

Holweg, M., Younger, R. & Wen, Y. (2022). The reputational risks of AI. *California Management Review,*

Humble, J. & Farley, D. (2010). *Continuous delivery: reliable software releases through build, test, and deployment automation*. Pearson Education, Hoboken.

Humble, J. & Molesky, J. (2011). Why enterprises must adopt DevOps to enable continuous delivery. *Cutter IT Journal, 24* (8), 6.

Jones, M. (2014). A matter of life and death. *MIS Quarterly, 38* (3), 895-A896.

Juergensen, T. (2000). *Continuous improvement: Mindsets, capability, process, tools and results*. Indianapolis: The Juergensen Consulting Group, I.

Jung, M., Niculita, O. & Skaf, Z. (2018). Comparison of different classification algorithms for fault detection and fault isolation in complex systems. *Procedia Manufacturing, 19*, 111-118.

Kane, G.C., Young, A.G., Majchrzak, A. & Ransbotham, S. (2021). Avoiding an oppressive future of machine learning: A design theory for emancipatory assistants. *MIS Quarterly, 45* (1), 371-396.

Kellogg, K.C., Valentine, M.A. & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals, 14* (1), 366-410.

Kim, G., Behr, K. & Spafford, K. (2014). *The phoenix project: A novel about IT, DevOps, and helping your business win*. IT Revolution Press, Portland, Oregon.

Knorr Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Harvard University Press, Cambridge, MA.

Kreuzberger, D., Kühl, N. & Hirschl, S. (2022). Machine Learning Operations (MLOps): Overview, Definition, and Architecture. *arXiv preprint arXiv:2205.02302,*

Krishnamurthy, R. & Yauch, C.A. (2007). Leagile manufacturing: a proposed corporate infrastructure. *International Journal of Operations & Production Management, 27* (6), 588-604.

Langley, A., Lindberg, K., Mørk, B.E., Nicolini, D., Raviola, E. & Walter, L. (2019). Boundary work among groups, occupations, and organizations: From cartography to process. *Academy of Management Annals, 13* (2), 704-736.

Latour, B. (2007). *Reassembling the social: An introduction to actor-network-theory*. Oxford University Press, Oxford.

Lawrence, T.B. & Dover, G. (2015). Place and institutional work: Creating housing for the hard-to-house. *Administrative Science Quarterly, 60* (3), 371-410.

Leavitt, K., Schabram, K., Hariharan, P. & Barnes, C.M. (2021). The Machine Hums! Addressing Ontological and Normative Concerns Regarding Machine Learning Applications in Organizational Scholarship. *Academy of Management Review,* (ja)

Lebovitz, S., Levina, N. & Lifshitz-Assaf, H. (2021). Is AI ground truth really "true"? The dangers of training and evaluating AI tools based on experts' know-what. *Management Information Systems Quarterly, 45* (3), 1501-1526.

Lebovitz, S., Lifshitz-Assaf, H. & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science, 33* (1), 126-148.

Leite, L., Rocha, C., Kon, F., Milojicic, D. & Meirelles, P. (2019). A survey of DevOps concepts and challenges. *ACM Computing Surveys (CSUR), 52* (6), 1-35.

Leonardi, P.M. (2011). When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. *MIS Quarterly, 35* (1), 147-167.

Leonardi, P.M. (2013). Theoretical foundations for the study of sociomateriality. *Information and Organization, 23* (2), 59-76.

Leonardi, P.M. & Barley, S.R. (2008). Materiality and change: Challenges to building better theory about technology and organizing. *Information and Organization, 18* (3), 159-176.

Leonardi, P.M., Nardi, B.A. & Kallinikos, J. (2012). *Materiality and organizing: Social interaction in a technological world*. Oxford University Press, Oxford.

Lou, B. & Wu, L. (2021). AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms. *MIS Quarterly, 45* (3), 1451-1482.

Mäkinen, S., Skogström, H., Laaksonen, E. & Mikkonen, T. (2021). *Who needs MLOps: What data scientists seek to accomplish and how can MLOps help?*, 2021 IEEE/ACM 1st Workshop on AI Engineering-Software Engineering for AI (WAIN). IEEE, pp. 109-112.

Martin, K. (2019). Designing ethical algorithms. *MIS Quarterly Executive June*, 129-142.

Mayer, A.-S., Strich, F. & Fiedler, M. (2020). Unintended Consequences of Introducing AI Systems for Decision Making. *MIS Quarterly Executive, 19* (4), 239-257.

Mazmanian, M., Orlikowski, W.J. & Yates, J. (2013). The autonomy paradox: The implications of mobile email devices for knowledge professionals. *Organization Science, 24* (5), 1337-1357.

McCulloch, W.S. & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics, 5*, 115-133.

Mishra, A. & Otaiwi, Z. (2020). DevOps and software quality: A systematic mapping. *Computer Science Review, 38*, 100308.

Mittelstadt, B.D., Allo, P., Taddeo, M., Wachter, S. & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society, 3* (2), 2053951716679679.

Muralidhar, N., Muthiah, S., Butler, P., Jain, M., Yu, Y., Burne, K., Li, W., Jones, D., Arunachalam, P. & McCormick, H.S. (2021). Using AntiPatterns to avoid MLOps Mistakes. *arXiv preprint arXiv:2107.00079,*

Ohno, T. & Bodek, N. (1988). *Toyota Production System*, 1 ed. CRC Press,

Orlikowski, W.J. (2005). Material works: Exploring the situated entanglement of technological performativity and human agency. *Scandinavian Journal of Information Systems, 17* (1), 6.

Orlikowski, W.J. (2007). Sociomaterial practices: Exploring technology at work. *Organization Studies, 28* (9), 1435-1448.

Orlikowski, W.J. (2010). The sociomateriality of organisational life: considering technology in management research. *Cambridge journal of economics, 34* (1), 125-141.

Orlikowski, W.J. & Scott, S.V. (2008). 10 sociomateriality: challenging the separation of technology, work and organization. *Academy of Management Annals, 2* (1), 433-474.

Orlikowski, W.J. & Scott, S.V. (2015a). The algorithm and the crowd. *MIS Quarterly, 39* (1), 201-216.

Orlikowski, W.J. & Scott, S.V. (2015b). Exploring material-discursive practices. *Journal of Management Studies, 52* (5), 697-705.

Pachidi, S., Berends, H., Faraj, S. & Huysman, M. (2021). Make way for the algorithms: Symbolic actions and change in a regime of knowing. *Organization Science, 32* (1), 18-41.

Pickering, A. (1995). *The mangle of practice [electronic resource] : time, agency, and science*. University of Chicago Press, Chicago.

Poppendieck, M. & Poppendieck, T. (2003). *Lean software development: an agile toolkit*. Addison-Wesley,

Raj, E., Buffoni, D., Westerlund, M. & Ahola, K. (2021). *Edge MLOps: An automation framework for aiot applications*, 2021 IEEE International Conference on Cloud Engineering (IC2E). IEEE, pp. 191-200.

Renggli, C., Rimanic, L., Gürel, N.M., Karlaš, B., Wu, W. & Zhang, C. (2021). A data quality-driven view of mlops. *arXiv preprint arXiv:2102.07750,*

Royce, W.W. (1987). *Managing the development of large software systems: concepts and techniques*, Proceedings of the 9th international conference on Software Engineering, pp. 328-338.

Ruf, P., Madan, M., Reich, C. & Ould-Abdeslam, D. (2021). Demystifying mlops and presenting a recipe for the selection of open-source tools. *Applied Sciences, 11* (19), 8861.

Rumelhart, D.E. & McClelland, J.L. (1986). *Parallel distributed processing : explorations in the microstructure of cognition*. MIT Press, Cambridge, Massachusetts.

Sako, M., Qian, M. & Attolini, J. (2022). Future of professional work: evidence from legal jobs in Britain and the United States. *Journal of Professions and Organization, 9* (2), 143-169.

Salama, K., Kazmierczak, J. & Schut, D. (2021). Practitioners Guide to MLOps: A Framework for Continuous Delivery and Automation of Machine Learning. *Google Could White paper,*

Sassen, S. (2006). *Territory, authority, rights : from medieval to global assemblages*. Princeton University Press, Princeton, N.J.

Schwaber, K. (1997). *Scrum development process*, Business Object Design and Implementation: OOPSLA'95 Workshop Proceedings 16 October 1995, Austin, Texas. Springer, pp. 117-134.

Schwaber, K. (2004). *Agile project management with Scrum*. Microsoft press,

Scott, S.V. & Orlikowski, W.J. (2014). Entanglements in practice. *MIS Quarterly, 38* (3), 873-894.

Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V. & Young, M. (2014). Machine learning: The high interest credit card of technical debt.

Senoner, J., Netland, T. & Feuerriegel, S. (2022). Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science, 68* (8), 5704-5723.

Slife, B.D. (2004). Taking practice seriously: Toward a relational ontology. *Journal of theoretical and philosophical psychology, 24* (2), 157.

Suddaby, R. (2014). *Editor's comments: Why theory?* Academy of Management Briarcliff Manor, NY, pp. 407-411.

Sugimori, Y., Kusunoki, K., Cho, F. & UCHIKAWA, S. (1977). Toyota production system and kanban system materialization of just-in-time and respect-for-human system. *The international journal of production research, 15* (6), 553-564.

Symeonidis, G., Nerantzis, E., Kazakis, A. & Papakostas, G.A. (2022). *MLOps-Definitions, Tools and Challenges*, 2022 IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC). IEEE, pp. 0453-0460.

Tamburri, D.A. (2020). *Sustainable mlops: Trends and challenges*, 2020 22nd international symposium on symbolic and numeric algorithms for scientific computing (SYNASC). IEEE, pp. 17-23.

Taylor, J.R. (2001). Toward a theory of imbrication and organizational communication. *The American Journal of Semiotics, 17* (2), 269-297.

Testi, M., Ballabio, M., Frontoni, E., Iannello, G., Moccia, S., Soda, P. & Vessio, G. (2022). MLOps: A Taxonomy and a Methodology. *IEEE Access,*

Treveil, M., Omont, N., Stenac, C., Lefevre, K., Phan, D., Zentici, J., Lavoillotte, A., Miyazaki, M. & Heidmann, L. (2020). *Introducing MLOps*. O'Reilly Media,

Tsamados, A., Aggarwal, N., Cowls, J., Morley, J., Roberts, H., Taddeo, M. & Floridi, L. (2021). The ethics of algorithms: key problems and solutions. *AI & SOCIETY,*

Vaast, E. & Walsham, G. (2005). Representations and actions: the transformation of work practices with IT use. *Information and Organization, 15* (1), 65-89.

Vartak, M. (2022). How to scale AI in your organization. *Harvard Business Review,*

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*

Vaughan, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. University of Chicago press,

Visengeriyeva, L., Kammer, A., Bär, I., Kniesz, A. & Plöd, M. Available at: https://ml-ops.org/content/mlops-principles (Accessed: August 17th 2021)

Waardenburg, L., Huysman, M. & Sergeeva, A.V. (2021). In the land of the blind, the one-eyed man is king: Knowledge brokerage in the age of learning algorithms. *Organization Science, 33* (1), 59-82.

Walsham, G. (2005). Agency theory: integration or a thousand flowers? *Scandinavian Journal of Information Systems, 17* (1), 11.

Yusuf, Y.Y., Sarhadi, M. & Gunasekaran, A. (1999). Agile manufacturing:: The drivers, concepts and attributes. *International Journal of Production Economics, 62* (1-2), 33-43.

Zammuto, R.F., Griffith, T.L., Majchrzak, A., Dougherty, D.J. & Faraj, S. (2007). Information technology and the changing fabric of organization. *Organization Science, 18* (5), 749-762.

Zbaracki, M.J. (1998). The rhetoric and reality of total quality management. *Administrative Science Quarterly, 43* (3), 602-636.

Zhang, Z., Hummel, J.T., Nandhakumar, J. & Waardenburg, L. (2020). Addressing the Key Challenges of Developing Machine Learning AI Systems for Knowledge-Intensive Work. *MIS Quarterly Executive, 19* (4), 221-238.

Zhou, Y., Yu, Y. & Ding, B. (2020). *Towards mlops: A case study of ml pipeline platform*, 2020 International conference on artificial intelligence and computer engineering (ICAICE). IEEE, pp. 494-500.

Chapter 6

Conclusion

6.1 Theoretical contributions

Drawing on investigations of how organisations use algorithmic technologies in practice, this thesis makes four main contributions. First, we contribute to research on the future of work, which has mainly focused on the labour displacement effects of algorithmic technologies (Acemoglu and Autor, 2011a; Autor, 2013, 2015; Autor et al., 2003; Frey and Osborne, 2017). Our results contrast with the view that algorithmic technologies will replace experts and professionals, and, instead, support more recent writing on potential work augmentation effects (Acemoglu and Restrepo, 2020; Sampson, 2021; Spring et al., 2022). We show that algorithms transform knowledge work by altering the structure of its constituting tasks. Because tasks within a job can each require different qualifications that affect its suitability for algorithmic technologies, they moderate potential productivity benefits and lead organisations to embrace hybrid work structures between algorithms and knowledge workers. Thus, confirming Sampson's (2021) work, our results suggest that research on the future of knowledge work should focus more on job redesign instead of job displacement.

Our second main contribution is also to the future work literature, which has traditionally focused on skills as the primary factor determining susceptibility to computerization (Acemoglu and Autor, 2011b; Autor, 2015; Autor et al., 2003; Frey and Osborne, 2017; Sampson, 2021). In contrast to the prevailing discourse, our research challenges the notion that creativity, social intelligence, and physical manipulation skills are the main determinants of enhancing knowledge work with algorithmic technologies. Instead, we posit that the key lies in an organisation's ability to capture, translate, and modify the underlying know-how of the practice. By doing so, we can unlock the potential for improving knowledge work with algorithmic technologies. Our findings further reveal that the opportunities for improvement are contingent upon the characteristics of the

process and the initial interpretations of algorithmic technologies by workers, who use algorithmic technologies as a way to achieve their goals, either by disrupting established structures or by maintaining their current status.

Our third main contribution is to the literature on process improvement in the context of knowledge work, which has highlighted that knowledge work exhibits contextual and structural characteristics, such as task ambiguity, that hamper improvement (Harvey et al., 2016; Spear and Bowen, 1999; Staats et al., 2011). We provide an alternative view in our study of the role of a deep learning-based decision models in improving inventory record inaccuracy – a classic problem in inventory control. We stress that algorithms can boost productivity in knowledge work by improving the representation of complex problems preceding operational decisions. Confirming the work of Bastani et al. (2020) and Mišić and Perakis (2020), our results show that generic machine learning algorithms are transforming the way operations management is practised. Specifically, our findings extend the work of Senoner et al. (2022), Ferreira et al. (2016), and Bernstein et al. (2019), who document that machine learning-based decision models can improve on quality management, demand forecasting, and revenue management; by highlighting that a deep learning-based decision model performs better than established inventory control methods; it creates a more accurate and fast representation of IRI, which led to swift and frequent problem-solving and improved the execution of subsequent replenishment operations.

Our fourth main contribution is to the literature on materiality and practice (Cecez-Kecmanovic et al., 2014; Jones, 2014; Orlikowski and Scott, 2008, 2015), which has primarily focused on examining how technology impacts organisations. We have highlighted the equally important notion that technology is shaped by organisations themselves. We showed that the development and operations practices of organisations

have a profound impact on the consequences of algorithmic technologies, thus emphasizing the generative potential of exploring the interplay between these two concepts. We further identified three relevant dimensions that can further advance our understanding of this interplay: the relationships between algorithm managers, developers, and users, the activities that algorithm developers and operators perform, and the data and automation artefacts that facilitate these practices. Our findings hold particular significance in the current context of debates surrounding the ontological and normative implications of algorithms(Leavitt et al., 2021; Lindebaum and Ashraf, 2021). Ultimately, this research can help inform the development of ethical and robust algorithmic technologies (Floridi et al., 2022; Martin, 2019; Mittelstadt et al., 2016; Tsamados et al., 2021).

By showing how an improved representation of external phenomenon leads to better decisions, we further show a somewhat more complex relationship between the drivers of task performance than suggested in traditional research on performance and expertise (Allen and Choudhury, 2021). Namely, we found that the ability to build accurate representations of external phenomena, not domain expertise, is the key antecedent to making operational decisions and improving efficiency. We found evidence, however, that the performance improvement brought by algorithmic decision models can be moderated by degradation due to errors creeping into the model, as well as employees' aversion to algorithm recommendations. These aspects were not part of our main investigation, yet warrant further research.

Finally, our findings highlight the complementary relationship between algorithmic technologies and process improvement practices. Research on the productivity effects of algorithmic technologies has argued that technology adoption must be accompanied by organisational changes and complementary investments to improve performance

(Brynjolfsson, 1993; Brynjolfsson et al., 2019). Our findings support and extend this view by showing which organisational changes help appropriate technology in knowledge work. We found that companies are more likely to realise the benefits of algorithmic technologies when they also deploy process improvement practices to re-engineer their processes, such as removing unnecessary steps, detaching redundant interdependencies, and streamlining communication. Other complementary practices include those that focus on creating early symbolic and substantive success to overcome resistance from stakeholders (Collins and Browning, 2019). In this sense, our results extend work on the applicability of process improvement practices in knowledge work contexts (Holweg et al., 2018; Radnor et al., 2012; Shah et al., 2008; Staats et al., 2011) by highlighting that some practices can deliver and sustain improvements in these contexts, in general, and help deliver algorithmic technologies, in particular.

6.2 Managerial implications

Our findings have several implications for managers seeking to redesign knowledge work with algorithms. A key insight is that strategies involving the automation of knowledge work should be implemented at the task and process level rather than the occupational level, as present studies commonly suggest (Frey and Osborne, 2017). To decide how to reallocate work between humans and algorithmic technologies, companies should differentiate between work that involves well-defined and emergent problems. The former fits into traditional schemata and are good candidates for algorithms, such as *"auditing the financial statements of a company"* (Autor, 2013). The latter requires interpreting novel and unfamiliar events to create collective agreement on what is going on and plan potential solutions, such as *"assessing the impact of an unfamiliar risk on the productivity of the company*" (Maitlis and Sonenshein, 2010; Weick, 1988, 1993). These tasks involving

emergent problems make poor candidates for algorithms; humans should perform them using intuition, collaboration, imagination, and abductive reasoning. We also highlight that the problem of automating work requiring creative skills is a technological one, and so, it may be possible to reallocate crate work to algorithms. But our findings suggest that success depends on the company's ability to first, adapt the process architecture in which the task is embedded, including removing unnecessary interdependences, streamlining communications, and reengineering the process to maintain quality; and second, to create symbolic success, such as convincing project sponsors, producing early benefits, and maintaining improvements over time. We further argue that the problem of automating work requiring physical manipulation skills is an economic one. While algorithms can be embedded into robots to perform activities that require manual dexterity and/or work in confined spaces, they may be too costly or risk the safety of the workers around them. Thus, it is more likely that work that requires physical manipulation skills will remain allocated to humans in the future.

Our results show that companies should avoid using algorithmic technologies for processes with low levels of standardisation. Particularly those processes that are tailored to customers' needs or that depend on experts' discretion tend to benefit less from automation algorithms. This observation also implies that organisations are more likely to succeed at automating/augmenting knowledge work when they re-engineer it, standardising as many steps and outcomes as possible. Based on our analyses, we further note that organisations benefit from automating work related to highly demanded company processes. For one, such highly repetitive processes allow the company to scale algorithmic solutions and compensate for the economic hurdles that may emerge from remapping the work of experts and professionals.

6.3 Limitations and further research

As with all case-based research, the generalizability of our findings is a major concern. We note that the cases presented in this thesis strived for naturalistic generalizability (Pratt et al., 2006), rather than statistical generalizability, and aimed at building insights from unique, rather than common, cases (Yin, 2017). Care must be taken when transferring our findings to other organisations. One key limitation stems from the budget and time constraints that lead us to focus on a specific part of the phenomenon of algorithms in knowledge work; We focused only on the characteristics of knowledge work that make it more amenable to algorithmic technologies. We thus did not examine other potentially influential factors, such as the reaction of knowledge workers to algorithms to resist improvement and automation initiatives. It may be that knowledge workers construct narratives about algorithmic technologies in strategic ways, protecting their agency and work, and ensuring control of key tasks; or it may be that experts' and professionals' aversion to algorithmic recommendations affect the performance of processes based on algorithmic technologies in the long run. Thus, two open questions are how do knowledge workers construct, resist and react to organisational efforts to redesign knowledge work with algorithmic technologies? And how do these social processes end up moderating the performance gains of algorithmic technologies?

A second limitation stems from our case study research design and access to data, particularly, for our qualitative study. Our results were shaped and limited by the people who we spoke with and observed, and by the information shared by the company. The participant companies may have strategically omitted information or prevented us from speaking with informants who had access to further insights. We nevertheless aimed to overcome this limitation by triangulating data with multiple external sources. Future

research may seek to include data from various stakeholders involved in the automation of knowledge work, for example, directors, suppliers, and customers.

A third limitation stems from our contexts, as our cases were built on engineering and retail firms, which likely shaped our findings. Knowledge workers in these sectors perhaps are more open to the adoption of new technologies, as they usually come from STEM backgrounds, and work in hierarchical organisations. And so, these knowledge workers are more likely to appropriate technology. It may be that knowledge workers at professional service firms, which are flatter and deployed knowledge workers from non-STEM backgrounds, use their power and react differently to these technological changes, structuring algorithmic work differently, and experiencing different effects on performance. Scholars could design future research to explore these differences. One area for future research may be to compare how knowledge workers at professional service firms adopt algorithmic technologies at work. This may include questions about the contingencies that affect the appropriation of benefits from algorithmic technologies (e.g., labour unionization status, firm size, firm age, organisation hierarchical structure, and degree of capital intensity among others).

In addition to the limitations of this study, we identify two critical themes for further research that could enrich the understanding of the impact of algorithmic technologies on knowledge work: (1) Knowledge destruction and transfer in the context of adopting algorithmic technologies, and (2) evaluating the impact of different types of algorithms on knowledge work improvement:

The first critical theme for further research revolves around the dynamics of knowledge destruction and transfer within organisations, particularly in the context of adopting algorithmic technologies. As we emphasized in Chapter 4, the integration of algorithmic technologies inevitably reshapes the organisation's knowledge regime. This

includes altering the knowing practices, the structural arrangements that govern knowledge, and the criteria used to assess the validity of knowledge outputs.

While such changes may enhance operational performance, they also lead to the destruction and transfer of knowledge. For instance, by automating the production of certain knowledge outputs, like creating auditing reports, an organisation might choose to remove auditors from the process. Over time, this leads to a lack of opportunities for auditors to learn from experience, and auditors over time may realise that they no longer know how to create auditing reports. Thus, introducing algorithmic technologies in knowledge work could lead to self-alienation, wherein workers perceive their knowledge as unfamiliar or alien. This process can trigger a transformative shift, causing workers to rediscover their identities within the organisation, in relation to the knowledge destructed by algorithms.

This may also pose challenges to knowledge transfer, as members of the organisation might no longer possess the necessary information and skills to solve emergent problems, nor can they effectively communicate these insights to others outside the organisation.

Consequently, further research should concentrate on comprehending the implications of adopting algorithmic technologies in knowledge work, particularly concerning knowledge destruction and transfer. Understanding how these changes impact the organisation's ability to share and maintain vital knowledge will be crucial for managing knowledge effectively in the era of technological advancement.

A third crucial area for future research revolves around evaluating the impact of various types of algorithms on knowledge work improvement. Throughout this thesis, we have focused on comprehending how algorithmic technologies, particularly machine learning (ML) and Robotic Process Automation (RPA) algorithms, influence knowledge

work. ML serves as an example of an inference-based algorithm, while RPA exemplifies a Logic-based algorithm. By investigating these instances, we aimed to develop a theoretical framework for understanding the impact of algorithmic technologies on knowledge work. However, we acknowledge that our findings may not be universally applicable to other algorithms, such as those traditionally used in Operations Research.

To advance our understanding, further studies can delve into whether the insights gained from this thesis are relevant to other algorithms or if there are notable differences. This exploration could lead to the development of an alternative taxonomy of algorithmic technologies. It is noteworthy that as machine learning architectures progress, particularly large language models, we witness increased variation in their capabilities, which warrants additional investigation. For instance, novel large language models display a level of reflexivity that bears semblance to human thinking, at least in a semiotic manner.

Hence, future research has the opportunity to refine our theory regarding how algorithmic technologies enhance knowledge work, considering the role of the algorithm type as a moderating factor. This investigation can help us better understand the nuances and potentials of various algorithms and their implications for knowledge work improvement. By exploring these aspects, we can make informed decisions about the adoption and implementation of algorithmic technologies in organisations, leading to more effective and efficient knowledge work practices.

References

Acemoglu, D. & Autor, D.H. (2011a). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics, 4*, 1043-1171.

Acemoglu, D. & Autor, D.H. (2011b). 'Skills, tasks and technologies: Implications for employment and earnings', in: Card, D., Ashenfelter, O. (Eds), *Handbook of labor economics*. Elsevier, Amsterdam, pp. 1043-1171.

Acemoglu, D. & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy, 128* (6), 2188-2244.

Allen, R. & Choudhury, P. (2021). Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion. *Organization Science, 33* (1), 149-169.

Autor, D.H. (2013). The "task approach" to labor markets: an overview. *Journal for labour market research, 46* (3), 185-199.

Autor, D.H. (2015). Why are there still so many jobs? The history and future of workplace automation. *The Journal of Economic Perspectives, 29* (3), 3-30.

Autor, D.H., Levy, F. & Murnane, R.J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics, 118* (4), 1279-1333.

Bastani, H., Zhang, D. & Zhang, H. (2020). Applied machine learning in operations management. *Available at SSRN 3736466,*

Bernstein, F., Modaresi, S. & Sauré, D. (2019). A dynamic clustering approach to data-driven assortment personalization. *Management Science, 65* (5), 2095-2115.

Brynjolfsson, E. (1993). The productivity paradox of information technology. *Communications of the ACM, 36* (12), 66-77.

Brynjolfsson, E., Rock, D. & Syverson, C. (2019). 'Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics', in: Agrawal, A. (Ed), *The Economics of Artificial Intelligence*. University of Chicago Press, Chicago, pp. 23-60.

Cecez-Kecmanovic, D., Galliers, R.D., Henfridsson, O., Newell, S. & Vidgen, R. (2014). The sociomateriality of information systems. *MIS Quarterly, 38* (3), 809-830.

Collins, S.T. & Browning, T.R. (2019). It worked there, so it should work here: Sustaining change while improving product development processes. *Journal of Operations Management, 65* (3), 216-241.

Ferreira, K.J., Lee, B.H.A. & Simchi-Levi, D. (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management, 18* (1), 69-88.

Floridi, L., Holweg, M., Taddeo, M., Amaya Silva, J., Mökander, J. & Wen, Y. (2022). capAI-A Procedure for Conducting Conformity Assessment of AI Systems in Line with the EU Artificial Intelligence Act. *Available at SSRN 4064091,*

Frey, C.B. & Osborne, M.A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114* (2017), 254-280.

Harvey, J., Heineke, J. & Lewis, M. (2016). Editorial for journal of operations management special issue on "professional service operations management (PSOM)". *Journal of Operations Management, 42-43* (1), 4-8.

Holweg, M., Davies, J., De Meyer, A., Lawson, B. & Schmenner, R.W. (2018). *Process theory: The principles of operations management*. Oxford University Press, Oxford.

Jones, M. (2014). A matter of life and death. *MIS Quarterly, 38* (3), 895-A896.

Leavitt, K., Schabram, K., Hariharan, P. & Barnes, C.M. (2021). The Machine Hums! Addressing Ontological and Normative Concerns Regarding Machine Learning Applications in Organizational Scholarship. *Academy of Management Review,* (ja)

Lindebaum, D. & Ashraf, M. (2021). The ghost in the machine, or the ghost in organizational theory? A complementary view on the use of machine learning. *Academy of Management Review,* (ja)

Maitlis, S. & Sonenshein, S. (2010). Sensemaking in crisis and change: Inspiration and insights from Weick (1988). *Journal of Management Studies, 47* (3), 551-580.

Martin, K. (2019). Ethical implications and accountability of algorithms. *Journal of Business Ethics, 160* (4), 835-850.

Mišić, V.V. & Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management, 22* (1), 158-169.

Mittelstadt, B.D., Allo, P., Taddeo, M., Wachter, S. & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society, 3* (2), 2053951716679679.

Orlikowski, W.J. & Scott, S.V. (2008). 10 sociomateriality: challenging the separation of technology, work and organization. *Academy of Management Annals, 2* (1), 433-474.

Orlikowski, W.J. & Scott, S.V. (2015). The algorithm and the crowd. *MIS Quarterly, 39* (1), 201-216.

Pratt, M.G., Rockmann, K.W. & Kaufmann, J.B. (2006). Constructing professional identity: The role of work and identity learning cycles in the customization of identity among medical residents. *Academy of Management Journal, 49* (2), 235-262.

Radnor, Z.J., Holweg, M. & Waring, J. (2012). Lean in healthcare: The unfilled promise? *Soc Sci Med, 74* (3), 364-371.

Sampson, S.E. (2021). A strategic framework for task automation in professional services. *Journal of Service Research, 24* (1), 122-140.

Senoner, J., Netland, T. & Feuerriegel, S. (2022). Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science, 68* (8), 5704-5723.

Shah, R., Goldstein, S.M., Unger, B.T. & Henry, T.D. (2008). Explaining anomalous high performance in a health care supply chain. *Decision Sciences, 39* (4), 759-789.

Spear, S. & Bowen, H.K. (1999). Decoding the DNA of the Toyota production system. *Harvard Business Review, 77* (5), 96-108.

Spring, M., Faulconbridge, J. & Sarwar, A. (2022). How information technology automates and augments processes: Insights from Artificial-Intelligence-based systems in professional service operations. *Journal of Operations Management, 68* (6-7), 592-618.

Staats, B.R., Brunner, D.J. & Upton, D.M. (2011). Lean principles, learning, and knowledge work: Evidence from a software services provider. *Journal of Operations Management, 29* (5), 376-390.

Tsamados, A., Aggarwal, N., Cowls, J., Morley, J., Roberts, H., Taddeo, M. & Floridi, L. (2021). The ethics of algorithms: key problems and solutions. *AI & SOCIETY,*

Weick, K.E. (1988). Enacted sensemaking in crisis situations. *Journal of Management Studies, 25* (4), 305-317.

Weick, K.E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly, 38* (4), 628-652.

Yin, R.K. (2017). *Case study research and applications: Design and methods.* Sage publications, Los Angeles.