



# Transitioning Towards a Proactive Practice: A Longitudinal Field Study on the Implementation of a ML System in Adult Social Care

Tyler Reinmund  
tyler.reinmund@jesus.ox.ac.uk  
University of Oxford  
Oxford, UK

Lars Kunze  
lars@robots.ox.ac.uk  
University of Oxford  
Oxford, UK

Marina Jirotko  
marina.jirotko@cs.ox.ac.uk  
University of Oxford  
Oxford, UK

## ABSTRACT

Politicians and care associations advocate for the use of machine learning (ML) systems to improve the delivery of adult social services. Yet, guidance on how to implement ML systems remains limited and research indicates that future implementation efforts are likely to encounter difficulties. We aim to enhance the understanding of ML system implementations by conducting a longitudinal field study with a team responsible for deploying a ML system within an adult social services department. The ML system implementation represented a cross-organisational effort to facilitate the department's transition to a proactive practice. Throughout this process, stakeholders adapted to numerous challenges in real-time. This study makes three contributions. First, we provide a description of how ML systems are implemented and highlight practical challenges. Second, we illustrate the utility of HCI knowledge in designing workflows for ML-assisted preventative care programmes. Finally, we provide recommendations for future deployments of ML systems in social care.

## CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Applied computing** → *Computing in government*.

## KEYWORDS

machine learning, implementation challenges, adult social care, field study

### ACM Reference Format:

Tyler Reinmund, Lars Kunze, and Marina Jirotko. 2024. Transitioning Towards a Proactive Practice: A Longitudinal Field Study on the Implementation of a ML System in Adult Social Care. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3613904.3642247>

## 1 INTRODUCTION

Adult social care (ASC) in England is weathering what many refer to as a “crisis” [2, 26]. The confluence of an ageing population, budget reductions in public services, and high attrition rates among

social care practitioners has led to a demand for care services that outweighs supply [33, 41]. In face of these operational and financial pressures, politicians and care associations advocate for the use of machine learning (ML) systems as a means to improve the efficiency and effectiveness of service delivery [69, 72, 102]. We define ML systems as software products that incorporate ML techniques to learn patterns from historical data to generate predictions on future instances. One application in ASC that garners particular attention is the use of ML systems to facilitate preventative care. These initiatives aim to prevent people from entering some state of acute need, whether it is experiencing a severe fall, social isolation, or admission to hospital [25, 29].

Despite the growing interest, guidance for social care organisations on how to implement ML systems remains limited [20] and the history of information systems in social care indicates that future implementation efforts are likely to encounter numerous human-computer interaction (HCI) related difficulties [12, 36, 50, 84]. And while HCI scholars have begun to research how social care practitioners use ML systems [17, 47], few, if any, studies have focused on the work required to put ML systems into practice – that is, how they are implemented. Therefore, we respond to calls in HCI to focus on the work of implementation [54] by gaining an understanding of the ML system implementation process in social care. We address the following research questions:

**RQ1:** How are ML systems implemented in social care organisations?

**RQ2:** What challenges arise during implementation?

**RQ3:** How do practitioners respond to these challenges?

Our team conducted a field study on the implementation of a ML system for a preventative care programme within a county council in England. County councils are local administrative bodies (i.e., “local authorities”) with responsibilities for public services such as education, transport, fire safety, and social care in a given region (i.e., a “county”). We conducted interviews with and observations of stakeholders from three participating organisations, supplemented with reviews of project documentation and artefacts. Our findings suggest that the ML system implementation represented a cross-organisational effort to facilitate the county council's transition to a proactive practice. Throughout this process, stakeholders experienced numerous challenges – many of which arise from the complex inter-organisational arrangements involved – which they adapted to in real-time by devising pragmatic resolutions.

This study makes three contributions. First, we provide a unique empirical description of how ML systems are implemented within social care organisations and highlight the practical challenges faced by stakeholders. In doing so, we show that focusing on implementation as an empirical site can support efforts in HCI to translate

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CHI '24, May 11–16, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0330-0/24/05

<https://doi.org/10.1145/3613904.3642247>

innovations into practical impact [54]. Second, we illustrate the utility of HCI knowledge in designing workflows for ML-assisted preventative care programmes. Third, we outline guidance for implementing ML systems in social care.

Before proceeding, it is important to clarify the aim of this study. The research team did not propose or design the ML system nor plan its implementation. Rather, we joined as observers to study how such decisions are made by practitioners in real-world situations. Our intention is not to present a case study on how ML systems *should* be implemented in social care, but instead to offer an open and critical elaboration on the complexities such projects present.

## 2 RELATED WORK & BACKGROUND

### 2.1 Information Systems Implementations in Social Care

Towards the end of the twentieth century, a series of governance and technological changes re-configured the social care landscape. The introduction of policy frameworks such as New Public Management and the coinciding development of information systems instigated a transition in social care organisations to prioritise practices of transparency, accountability, and standardisation [12, 34, 76]. Spurred by these developments – as well as a series of high-profile implementation failures [e.g., 84] – scholars in the Social Work literature began to focus on how these systems affected social work practice, the factors that shape their use, and techniques for engaging social workers in the design process [76, 108].

Written primarily from the perspective of social workers, one strand of scholarship questions the objectives underlying the design and implementation of information systems. Converging on a similar critique, these scholars contend that such information systems are designed with managerial goals in mind to the neglect of front-line workers [12, 36, 76]. In a retrospective focus group discussion on the failures of the Integrated Children's System in England, Wastell & White [98] argue that the system was designed primarily as a mechanism for auditing, thereby requiring social workers to complete complex forms and standardising their workflows to “squeeze out” social workers' discretion. Elaborating the implications of this change, Parton [76] writes that these practices may “become, in theory, more transparent and accountable,” but that these supposed benefits are counterbalanced by the loss of discretion for the social worker and the disappearance of “knowledge which cannot be squeezed into the required format.”

Much research has focused on the issues facing the design of information systems in social care; as the list below highlights, these challenges resonate with HCI discussions on technology in other domains of care [cf. 107].

- **Creation of administrative burden:** The completion of forms and inputting of data into information systems takes up time that social care practitioners feel could be better spent serving their clients [12, 36, 50, 84].
- **Poor workflow integration:** The systems regiment workflows that do not support the contingent and collaborative reality of social care practice [36, 93, 98].
- **Insensitivity to context:** Social care practitioners perceive that systems do not adequately represent disability and social

welfare knowledge [36, 84] and limit their ability to construct detailed narratives of individual clients [84, 98].

- **Exclusion from design process:** Social care practitioners are not given the opportunity to contribute to decisions during design, implementation, and system administration [12].

Despite these less favourable readings, several studies note how some social care practitioners retain positive expectations and perceptions of information systems in terms of their potential to benefit their practice [36, 84]. We chart these limitations as a basis for improving design initiatives in ML. Specifically, we employ these insights in Section 5 to conceptualise alternative workflows for ML systems that support service user engagement and satisfy the interaction needs of social care practitioners.

### 2.2 HCI Research on Machine Learning for Social Care

In the ensuing decades, technological developments in social care transitioned from the collection and storage of data to increasingly sophisticated means of its analysis [35, 37]. Garnering the attention of Social Work and HCI scholars alike, this phenomenon has inspired reflections on the opportunities data science presents for social care [14, 21], critical reflections on its challenges [35, 37], and a growing body of empirical work that explores how stakeholders perceive and experience ML systems [3, 11, 48, 86].

How stakeholders perceive the use of ML systems in social care is not uniform: it varies across people, use cases, and over time as they have recurring interactions with the technology. For example, Brown et al. [11] conduct a series of participatory design workshops to understand the perspectives of those who are impacted by ML systems in the US child welfare system. By bringing together families, front-line workers, and other practitioners, the authors identify a crucial aspect of trust that highlights the limitations of a perspective that emphasises technological design over other modes of intervention: distrust in the US child welfare system at large conjures distrust in their use of ML systems. Stapleton et al. [86] adopt a similar approach but begin by asking whether predictive risk models should be used in child protective services at all; their results point to a deep-seated concern that the use of such technologies perpetuate existing problems of punishment, disempowerment, and racial bias in child welfare.

These findings do not imply that interest in applying ML systems to support social care organisations should be abandoned. Several papers suggest that front-line workers see merit in applying the technology towards other use cases, such as the improvement of organisational processes or for evidential support to justify their own decisions [3, 44]. Kawakami et al. [48] summarise this reconceptualisation as a transition from “static predictions” to “communication tools.”

Recently, HCI scholars have turned towards understanding how ML systems impact work practices within social care. In a study with the Allegheny County Office of Children, Youth, and Families (CYF), Kawakami et al. [47] show how case workers draw upon contextual information not captured in the ML system to calibrate their reliance on its output, employ informal strategies to learn about its behaviour, and are influenced by organisational pressures

and incentives to use the system. Similarly engaging with the CYF, Cheng et al. [17] analyse four years of data on referral decisions and conduct contextual inquiries with case workers and supervisors to explore how practitioners reduce racial disparities in algorithmic decision-making.

Our study builds upon this work in several ways. First, we move away from Allegheny County and focus on a novel empirical site; this shift allows us to reveal how the economic and operational states of adult social services departments in England have tangible implications for how ML systems are designed and deployed, manifest through various implementation challenges. In turn, we contribute to past work that has shown how broader structural features in contexts of use affect stakeholder experiences with ML (and less sophisticated algorithmic) systems, such as resources, governance mechanisms, workplace procedures, and organisational incentives [47, 82]. To these considerations, we draw on our longitudinal field study to emphasise the role of business models, the restructuring of public services, and distributions of expertise. Finally, recent work focuses on how practitioners use ML systems within their existing work practices [17, 47, 48]. We instead focus on a practice **in the making**, revealing the joint development of a technical system and preventative care workflow and the influence of implementation decisions on matters of interaction experience.

### 2.3 Adult Social Care in England

ASC in England encapsulates a range of services delivered by numerous organisations and individuals, and touches upon the lives of an estimated 10 million care recipients, social care practitioners, and informal carers [70]. Adults with a physical disability, learning disability, or physical and mental illness receive assistance with essential daily activities such as eating, washing, and socialising [19]. This care is provided by a complex system of local authorities, health, housing, and welfare services, volunteer organisations, and informal carers, such as family and friends [19, 89].

In contrast to the country's health care system which is free at the point of use, formal ASC services are paid for either by the local authority or individually; to be eligible for publicly-funded care, service users undergo needs and means testing [103]. England's central government is responsible for setting policy, enacting legislation, and distributing funds to local authorities. Meanwhile, local authorities have statutory responsibilities to assess individuals' care needs and financial status and commission care services from private providers [19, 39].

Despite its scale and centrality to the lives of many, ASC has suffered from sustained under-investment and struggled to keep up with increasing demand [13, 19, 83, 87]. Researchers and social care associations have written extensively on the many factors that have led to this condition, spanning from geographic to demographic, financial to political. For example, Hamblin [41] and Wright [102] describe an ageing population, especially among rural and coastal areas, with increasing co-morbid and chronic conditions, high attrition rates for underpaid social care workers, and budgetary cuts for local authorities as leading to a situation where demand for social care services far outpaces supply. Awareness of these challenges is not new: successive governments have promised to reform the sector in one way or another since the 1990s [33].

In response to these challenges, policy directives and industry initiatives have forwarded digital technology and, more recently, ML as viable solutions [31, 69, 72]. Such proposals have been met with alacrity among practitioners who seek to leverage ML for care-related decisions [10, 102]. One such area that has been targeted by scholars and practitioners alike is falls detection. Each year, around 30% of older adults living at home fall at least once, and these incidents impact an individual's quality of life and health, and introduce costs to the health and social care system [66]. There has been substantial interest in designing interventions to predict an individual's likelihood to fall, employing technologies such as wearable devices and at-home sensors, for example [27].

While numerous policy and research reports suggest that such capabilities facilitate preventative and personalised service provisioning, thereby reducing costs and improving service quality [8, 70, 104], the evidence base to support such claims remains nascent [10, 22, 55]. Further, several reports go on to describe features of ASC that may inhibit the use of ML, such as low levels of data science competencies, under-developed technical infrastructures [8, 20, 72], and unrepresentative, incomplete, and infrequently collected data sets [16, 32, 37].

## 3 FIELD STUDY

### 3.1 Research Site

To answer our research questions, we conducted a longitudinal field study on the implementation of a ML system (the "project") from December 2022 to May 2023. The ML system was positioned as an intervention to help identify older adults who are at risk of experiencing a fall. Adults identified to be at risk of falling are deemed eligible for inclusion in a preventative care workflow. The ML system and preventative care workflow are described in Section 3.2.

This project was conducted by three geographically dispersed organisations: a county council, consultancy, and technology provider. The local authority is a county council in the East of England with both rural and urban districts. The county in which the implementation was conducted has a two-tier local government organisation: responsibility for public services in the county are distributed across a county council and several district councils. For example, services such as adult and children's social care are provided by the county council, whereas responsibility for housing and benefits services are spread amongst the various district councils. The consultancy is headquartered in South East England; yet, due to the firm's client-facing work, consultants are spread throughout the country, working from client sites, co-working spaces, or their own homes. The consultancy proposed the ML system, and managed and contributed to the design of the system and workflow. Meanwhile, the technology provider is headquartered in South East England and provides a data aggregation platform for local authorities.

While we only focused on one research site, we believe this setting to be representative of ML projects in social care organisations in England. Most local authorities utilise third party organisations to procure and implement digital systems [8], including those that incorporate ML [29]. Similarly, the challenge of addressing rising demand amid resource constraints facing the county council is experienced throughout the country. Additionally, by providing

detailed descriptions of the research context throughout the paper, we facilitate comparison of our study with related qualitative research to support the transferability of findings. Finally, other HCI researchers have successfully utilised single sites for related work [e.g., 44, 47, 61, 62, 82, 109], and we draw on their experiences to inform how we represent our findings.

## 3.2 Falls Prevention Programme

The ML system studied in this paper is composed of multiple bespoke and off-the-shelf software products. These components include the systems that collect, store, and pseudonymise the input data, the natural language processing (NLP) component that identifies risks within case notes, the ML model for falls prediction, and the data pack which serves as the user interface for the system predictions, illustrated in Figure 1.

As an input, the ML system analyses care package information and case notes stored in the adult social service department's **client management system** (CMS). Case notes are unstructured text data that social care practitioners create during their interactions with service users. The NLP component – known as the **risk framework** – extracts risks from the case notes by identifying keywords. This NLP analysis generates the **master risk table**: a structured data set that presents occurrences of risks for each service user. Using different combinations of risks as features, a binary classification algorithm was trained on this structured data set to create a **ML model** to predict whether a service user is likely to have a fall risk appear in their case notes in the next 9 months. The ML model calculates a probability score for each service user, and those with a score over a set threshold are identified to be at risk of falling. Finally, the predictions generated by the ML model are merged with demographic data and transferred to a **data pack**: an *Excel* workbook that serves as the interface through which stakeholders interact with the ML model's predictions.

The preventative care workflow is depicted in Figure 2. Once the ML system identifies service users who are at risk of falling, they are allocated to the county council's new falls prevention program. After a series of preliminary checks to ensure that non-eligible individuals are not contacted – which include identifying and excluding those who are deceased, in residential care, or under the age of 65 – the list of service users are shared with the Population Health Contact Centre. Members of this team are non-clinical practitioners: while they do not have a background in health or social care, they have experience performing a similar role for other public health initiatives. This team contacts each service user to conduct an over-the-phone assessment, asking a series of predefined questions; the purpose of these **holistic conversations** is to gain a better understanding of the factors that lead to an individual being at risk of falling. Based on this discussion, the call centre team member allocates the individual to an appropriate prevention service, such as exercise classes for mobility and strength or a home safety assessment.

## 3.3 Data Collection

From December 2022 to May 2023, the lead author joined the consultancy team as a participant observer. Throughout the field study, the lead author utilised three forms of data collection: participant

observation, interviews, and document and artefact analysis. Implementations occur over long periods of time; therefore, engaging with participants throughout the duration of the project was essential. Further, participant observation allows for a deep understanding of social interaction in a specific context and allows the research team to explore discrepancies in participants' verbal accounts [6, 23]. For example, Beede et al. [7] successfully employed participant observation to unearth the socio-environmental factors that affected the use of a deep learning system for diabetic retinopathy screening. Finally, the use of three methods supports the validity of findings through triangulation [23].

**3.3.1 Participant Observation.** The lead author joined the consultancy team as an “active member” [24]. This involved working on-site with stakeholders – at the county council's office or the consultancy's co-working space – and participation in various virtual and in-person meetings.

Inclusion of stakeholder groups varied across the meetings. Overall, discussions on the design of the ML system were restricted to the consultancy and technology provider, while the design of the workflow involved stakeholders from the county council and consultancy. We describe the implications of this differential participation in Section 4.1.3. Each of the recurring meetings took place over a virtual conferencing platform and was facilitated by one of the consultants through the use of project planning and collaboration software, such as *Asana*,<sup>1</sup> *Miro*,<sup>2</sup> and *PowerPoint*.<sup>3</sup>

To protect stakeholder confidentiality and ensure open discussion, meetings were not audio recorded. Instead, the lead author recorded handwritten notes on individuals' phrases, conversations, and expressions, and the environment in which the meeting took place. Immediately after, he expanded on the initial notes by filling in additional detail that was not captured instantaneously. At the end of each day, these notes were typed into *NVivo* 12<sup>4</sup> for analysis. The primary meetings and workshops the lead author attended are outlined in Table 1. The list is not comprehensive; ad-hoc meetings arose throughout the fieldwork.

**3.3.2 Interviews.** To extend the findings from the observations, the lead author conducted 13 formal semi-structured interviews and 40+ informal interviews with members of the project team. Thirteen semi-structured interviews were conducted with 4 practitioners from the county council, 1 from the technology provider, and 8 from the consultancy from varying levels of seniority. Each stakeholder was interviewed individually on a video conferencing platform. The interviews ranged from 35 to 60 minutes. Each interview focused on stakeholders' roles and involvement on the project, motivations behind the development of the ML system, perceived challenges, and perceptions of the ML system.

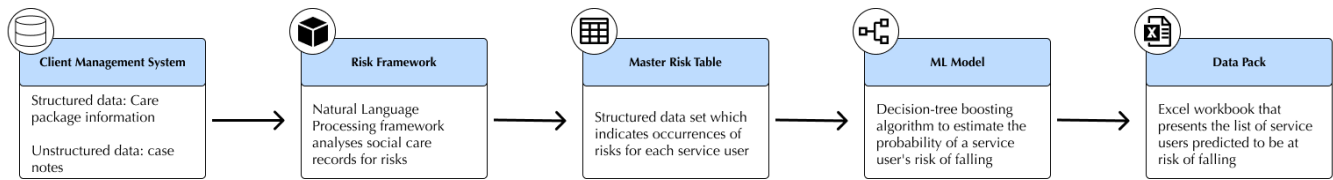
Throughout this period, the lead author also held 40+ informal interviews with members of the project team to address questions as they arose and validate observations, such as project challenges and experiences with collaborators. These conversations were held both in-person and virtually over a video conferencing platform.

<sup>1</sup><https://asana.com/>

<sup>2</sup><https://miro.com>

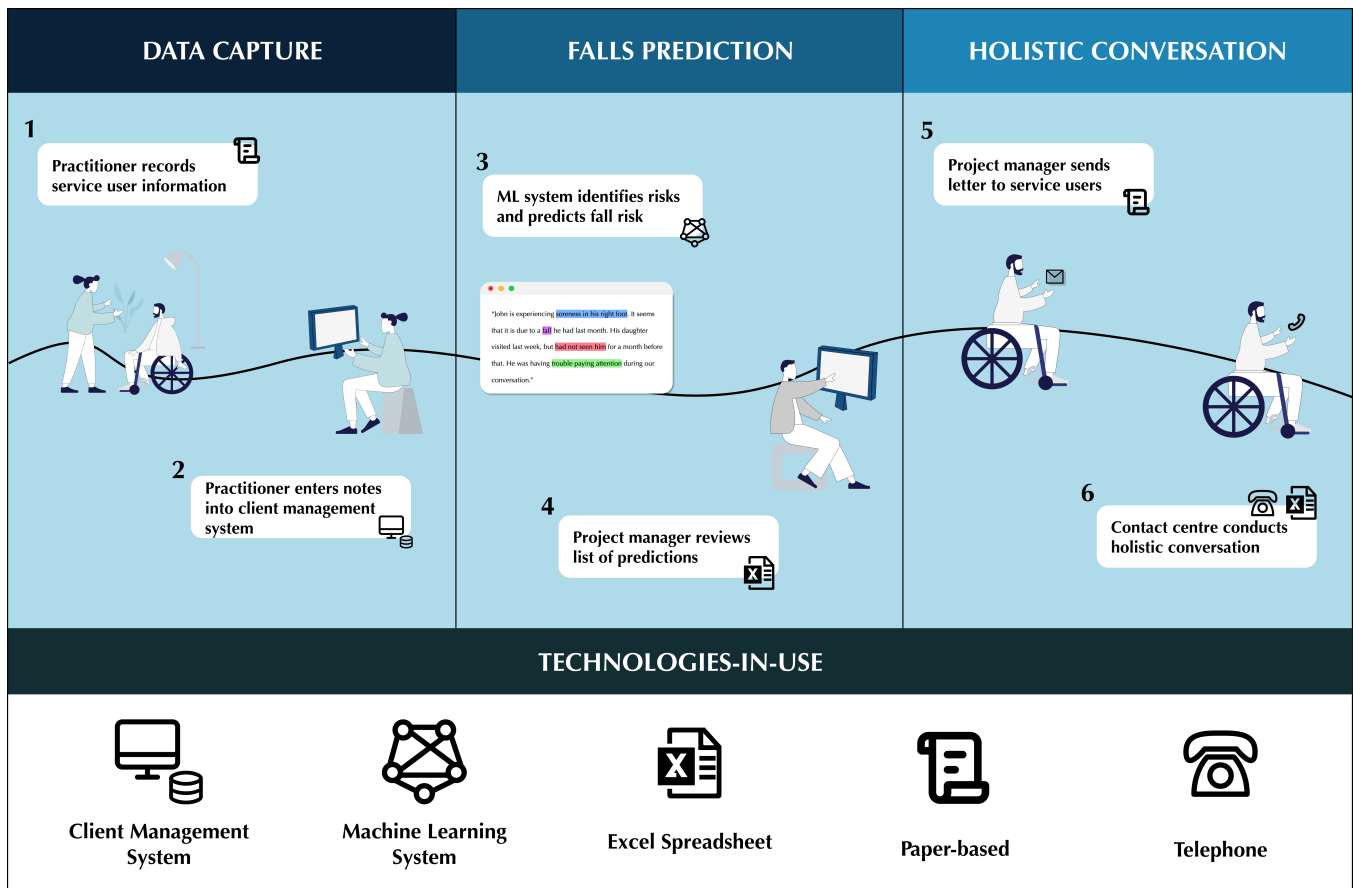
<sup>3</sup><https://www.microsoft.com/en-gb/microsoft-365/powerpoint>

<sup>4</sup><https://lumivero.com/products/nvivo/>



**Figure 1: A diagram of the ML system.** Care package information and case notes are stored in the client management system. Pseudonymised data is analysed by the risk framework to generate the master risk table, a structured data set of occurrences of risks by service user. This data set is used to train a ML model to predict each service user’s risk of falling. The list of predictions is disseminated in an Excel workbook referred to as the data pack.

## CURRENT WORKFLOW



**Figure 2: The workflow spans three phases, beginning with capturing data as case notes, generating predictions from case notes, and conducting holistic conversations with service users.** Most steps are facilitated by the use of some technology, spanning in complexity from paper-based letters to ML systems.

Interviews were not audio recorded. This decision was made for two reasons. During initial conversations with stakeholders, many expressed a concern with having their voice recorded, despite the research team’s measures to protect anonymity. Second, much of the work done in the public sector is sensitive, and there is a concern that research may lead to reputational damage for the organisation. Inspired by the approach of Veale, van Kleek, & Binns [92],

the lead author took copious notes during interviews to capture verbatim phrases, expressions, and sentences when possible. Notes were typed up immediately afterwards to ensure that details not captured instantaneously were retained; identifiable information was pseudonymised during this process.

**3.3.3 Documentation and Artefacts.** Finally, we collected documentation and artefacts produced throughout the project. These sources

**Table 1: Overview on primary meetings and workshops attended as part of participant observation. This is not a comprehensive representation of all forms of participation as ad-hoc meetings arose throughout the fieldwork. ■ marks meetings in which the organisation was involved, and □ marks those who were not present.**

Meeting / Workshop	Cadence		Organisation		
			County Council	Consultancy	Technology Provider
Project Kick-off	20 hours	One-off	□	■	□
Development Kick-off	0.5 hours	1 x wk.	■	■	■
Development Stand-up	0.5 hours	3 x wk.	■	■	■
Sprint Retrospective	1 hour	1 x mo.	■	■	■
Implementation Group Weekly	1 hour	1 x wk.	■	■	■
Implementation Group Workshop	2 hours	One-off	■	■	□
Project Review	7 hours	One-off	□	■	□

included *PowerPoint* presentations and *Word*<sup>5</sup> documents developed by stakeholders, recorded vendor tutorials, external training material for social care practitioners, and anonymised data analysis outputs from *Excel*<sup>6</sup> workbooks and *SQL* queries. Presentations, video recordings, and data analysis outputs created by stakeholders confirmed observations and provided insight into the motivations and expectations underlying the project.

### 3.4 Data Analysis

The lead author conducted data analysis as per the approach outlined by Miles, Huberman, & Saldaña [59]. This qualitative data analysis approach consists of two cycles. In the first cycle, the lead author assigned concept, process, and *in vivo* codes to data: concept codes represent higher-level abstractions, process codes depict conceptual action in data and are represented by a verb, and *in vivo* codes capture stakeholder language and are indicated by quotation marks.

During the second cycle, first cycle codes were written on a note card, alongside their type and a preliminary definition. These codes were then organised into patterns based on similarities. Once all notes were grouped, the patterns were recorded in an open-source concept mapping tool to preserve the analysis. The patterns generated during this cycle served as the basis for Section 4.

### 3.5 Researcher Role & Positionality

As noted before, the lead author adopted the role of an active member in the group under study. The reason behind this decision was to improve the reciprocity of the research relationship: rather than solely observing action, the lead author supported certain tasks for the consultancy when requested. Specifically, the lead author supported the consultancy by researching the medical literature on effective falls prevention interventions, determining an appropriate sample size, and identifying valid evaluation measures. Additionally, the lead author presented the emerging findings from the fieldwork to leadership within the consultancy and county council.

Each member of our research team has a unique background that influences our perspectives. We come from a range of academic disciplines, such as HCI, science and technology studies, artificial

intelligence, and requirements engineering. We are all researchers based in England with experiences studying diverse phenomena from work practices in financial trading rooms to interactions with social robots. Additionally, the lead author has a professional background in consulting for local government – albeit for a different organisation in another country – and academic training in HCI and machine learning that supported his integration into the research site.

### 3.6 Informed Consent & Participant Anonymity

Several steps were taken to ensure that stakeholders understood their involvement in the research and to protect their anonymity. Prior to entering the research site, the lead author's role was communicated to the project team following approval from leadership in the county council. Additionally, whenever encountering a new stakeholder, the lead author informed him or her about his role as a researcher, described the objective of the research, and requested consent. This was not possible in all situations – such as in meetings with a large number of attendees. In these instances, the lead author did not collect specific data on individuals for whom he was unable to inform. Additionally, an information sheet and consent form was shared when contacting individuals for interviews.

To ensure that stakeholders are not identifiable to others in their respective workplaces, we employ pseudonyms. Individuals from the county council, consultancy, and technology provider are indicated by **CC1–CC18**, **C1–C9**, and **TP1–TP2**, respectively. We present high-level participant information in Table 2; we do not provide a direct association of pseudonyms to roles to prevent the re-identification of quotes by colleagues. Quotes extracted from internal reports or presentations are indicated by **D**. Some sensitive quotes are not attributed a pseudonym to further protect the confidentiality of our stakeholders, especially given the small number of individuals on the project team. To support non-identifiable distinctions between the actions of members of each stakeholder group, individuals within the county council are referred to as practitioners and the consultancy as consultants.

This research received ethical approval by the University of Oxford's Computer Science Departmental Research Ethics Committee (DREC) and the county council's research governance panel.

<sup>5</sup><https://www.microsoft.com/en-gb/microsoft-365/word>

<sup>6</sup><https://www.microsoft.com/en-gb/microsoft-365/excel>

**Table 2: High-level information of stakeholders engaged with during the field study. In total, we interacted with 29 stakeholders from a diversity of roles across three organisations.**

Organisation	Roles	Count
County Council	Occupational Therapist	1
	Programme Manager	6
	Assistant Director	4
	Project Manager	1
	Analyst	1
	Specialist	5
Consultancy	Consultant	3
	Senior Consultant	2
	Manager	1
	Director	1
	Data Scientist	2
Technology Provider	Project Manager	1
	Data Scientist	1

## 4 FINDINGS FROM FIELD STUDY

### 4.1 Transitioning Towards a Proactive Practice

Project stakeholders conceptualised the ML system implementation as an effort to facilitate a transition within the county council from a “*reactive service to an increasingly proactive, preventative one*” [C8]. To establish proactive “*ways of working*” [C2, C7] and “*practices*” [CC8], a complex network of stakeholders held together by a multiplicity of motivations collaborated to integrate the output of a ML system into a novel preventative care workflow instituted through a pilot implementation design. As this section shows, the work required to implement an ML system is highly interdisciplinary, spanning traditional organisational boundaries, and entails the integration of people, organisational goals and processes, and legacy and novel technical systems.

**4.1.1 Complex Network of Stakeholders.** The project brought together a network of cross-sector groups, many of which had no prior experience working with one another. Each contributed to different phases of the project, from development of the ML system to delivery of the preventative care services, as depicted in Figure 3. The relationships between these organisations – in particular, the county council and the consultancy, and the consultancy and the technology provider – emerged from perceived gaps in their own and each other’s capabilities.

Much of the work was structured around the involvement of three primary organisations: the county council, the consultancy, and the technology provider. Stakeholders from each organisation formed a group responsible for delivering the project (the “project team”). Membership on this team was dynamic: different stakeholders departed and entered the project team over the course of the project.

Most of these groups had not previously worked with one another; departments within the county council had no experience collaborating with each other, nor did different organisations across the health and care sectors. As a result, this project led to the

establishment of relationships across traditional inter- and intra-organisational boundaries.

Driving much of the implementation was the consultancy which had been contracted by the county council. As explained by several of the consultants, the consultancy was approached by the county council to improve the referral process at the “*front door*” of the adult social services department [C4, C8]. In ASC, the front door refers to the initial point of contact between an individual – either someone looking for care themselves or for another person – and a county council’s ASC services. During the early stages of this broader programme, members of the consultancy reviewed cases that progressed to adult social services’ front door, leading to the conclusion that a substantial number of cases could have been prevented [C5, C8, D]. The consultancy used this conclusion as one basis to propose the creation of a preventative care programme – composed of a ML system and accompanying workflow – to proactively identify older adults at risk of escalating needs.

A core and novel component of this proposal, from the perspective of both the county council and consultancy, was the use of the ML system. Given that both organisations had limited experience developing ML systems, the consultancy enlisted a technology provider as a subcontractor. This relationship stemmed from a confluence of differences in expertise, professional relationships, and the state of the social care technology market: the provider’s ability to compensate for the consultancy’s lack of data engineering capabilities [C8, C9], a pre-existing relationship between the consultancy and the technology provider [TP1, C3, C8], and a dearth of competitors [C3, C8].

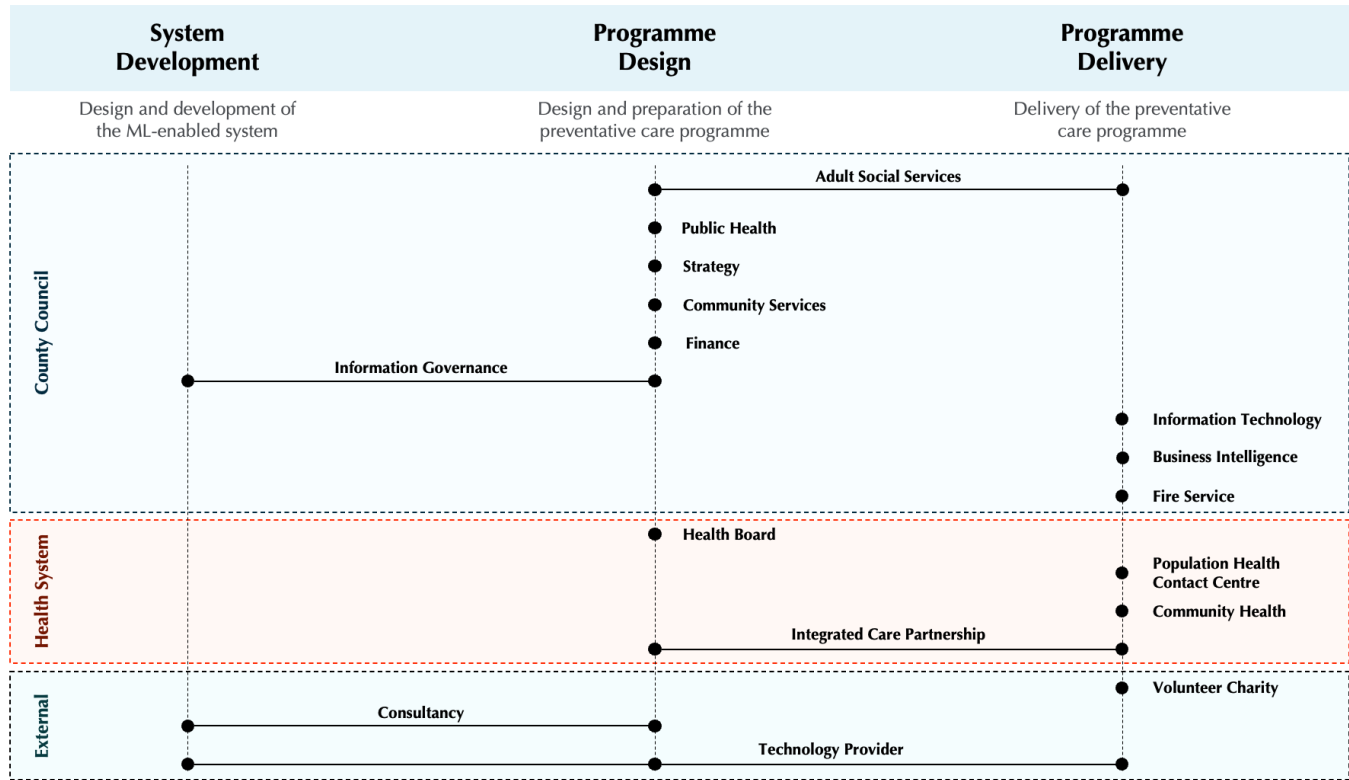
**4.1.2 Motivations for Developing a ML System.** Stakeholders provided varied responses when explaining the motivations that drove the decision to develop the ML system. Broadly, each explanation falls into one of four types: economic, operational, welfare, and strategic. We describe each motivation type in the subsequent paragraphs. The first two explanations were expressed by all stakeholders; in comparison, those that surface strategic reasons were shared by senior members of the consultancy who were involved in proposing the project to the county council.

The economic explanation characterises the decision to focus specifically on falls as a targeted response to an identified problem afflicting the county council. As a consultant describes:

**C5:** “*I believe there was some analysis done, and some scary statistic about the number of people experiencing a fall – like a third of people over the age of 65 – and that there is a higher chance that those people will need a package of care after the fall. And the cost associated with that care package. The idea was that if we could prevent those falls from happening, that would lead to a significant cost savings.*”

Related to the economic motivations spurring the project, others highlighted the operational challenges facing the county council and the welfare of its service users. At one end, there is increasing demand due to an ageing population, with “*queues that are already out the door*” [CC5]. And at the other, the county council struggles with limited capacity due to low levels of resources and difficulties recruiting and retaining staff [C4, CC1, CC2, CC7, CC8]. As such, these stakeholders hoped that preventing falls will lead to a





**Figure 3: The complex network of stakeholders involved in the transition towards a proactive practice. Disparate departments from the county council, organisations in the health system, and private and volunteer sector organisations occupied various roles across development of the ML system, and the design and delivery of the preventative care programme.**

reduction in demand for care services [CC8, C4] and improve the well-being of service users [D, CC8]. While these are important operational, welfare, and economic reasons, statements made by other members of the consultancy challenge the representation of this decision as one driven purely by analysis.

For the consultancy, the decision to develop a ML system for falls prevention emerged as part of an incipient market strategy. Members of the consultancy explained that there is a “gap in the market” [C2] associated with adopting preventative care programmes. In turn, the consultancy is developing a market strategy that focuses on using ML systems to instigate organisational change within local authorities [C2, C3]. Members of the consultancy claimed that there is increasing demand amongst local authorities to deliver preventative care programmes [C3, C4, C8]. Additionally, consultants perceived that the financial benefits associated with reducing falls would realise in a relatively shorter timeline as compared to other types of prevention programmes [C8]; therefore, the consultants hoped this project would provide a case study with quantifiable benefits to help extend this service to both different localities as well as other social and physical infirmities [C7, C8]. One consultant described the excitement among the civil servants and politicians “about taking this to health, all the remaining districts, and even the police” [C8] while practitioners propose that “this is not just going to look at falls, but also things like social isolation” [CC2].

**4.1.3 System and Workflow Design Approach.** Development of the ML system and the preventative care workflow were treated as two separate elements, performed by non-overlapping groups of stakeholders. ML system development followed a model-centric approach, while workflow design focused on interactions between the county council and service users. A consequence of these orientations was that the eventual primary user of the ML system was never defined.

ML system development preceded the definition of the workflow and emphasised optimising predictive performance. This model-centric approach entailed the prioritisation of precision as the main evaluative criterion for design decisions: how to specify the target variable and which features to include were driven by whatever led to a measurable improvement in precision [C2, TP2].

This work was performed exclusively by data scientists from the consultancy and technology provider. One consultant reflected on this approach, recognising that “I know that’s not the right [approach], but that’s the answer. It was mostly because of time pressures around the project” [C8]. Others within the consultancy and county council similarly attributed the lack of involvement of practitioners to time constraints among staff and the project’s short timeline [CC2, CC8, C6]. A consequence of the lack of involvement was that stakeholders had variable understandings of the ML system. Specifically, those who were not involved in its development, such



as practitioners, held a high-level understanding of the system; their knowledge was restricted to what was communicated to them in presentations by the technology provider or the consultancy [C8, CC2, CC5, CC8]. Whether practitioners desired to be involved in system design remains ambiguous: some expressed concern that their lack of technical expertise would preclude them from participating [CC2, CC5], while another recognised that a lack of involvement led to confusion over how the ML system functions [CC8].

Separate from development of the ML system, a group of practitioners from the county council and a consultant focused on designing the workflow. This group met each week over a virtual conferencing platform to discuss several aspects of the workflow, most of which centred around the interactions between the county council and service users: writing the letter to initially contact service users, defining a script for the Population Health Contact Centre, and determining which services to offer through the programme.

As most attention was being directed towards improving the predictive performance of the ML system or elaborating the interactions between service user and county council, one stakeholder group was neglected: the actual user of the ML system. The following two excerpts illustrate how the question of “who is the user?” remained unanswered across several months, even until after the system was deployed:

**TP2:** “It’s a thing I’m a little unsure on is how they will use the output [of the ML system]?”

**C7:** “Yeah that’s something we need to speak to [business intelligence team] and my manager about.” [Field notes: 24 January 2023]

**Researcher:** “Who is the primary user of the system?”

**CC7:** “As in a person? Hmm. that’s a good question. I’d almost say it’s [technology provider] but is that right? Well we now have [project manager] but I’m not sure about that. So I don’t know who would have a shortcut on their computer for the model and could open it up and do the modelling.” [Interview: 30 May 2023]

**4.1.4 Proving Effectiveness Through a Pilot Design.** To demonstrate that the preventative care programme can achieve the project’s operational, economic, and welfare objectives, the project adopted a pilot implementation design that ascribed significant importance to the role of evaluation.

There are several reasons as to why programme evaluation occupied such a prominent place within the project. Primarily, the consultancy positioned the programme as an opportunity for the county council to realise cost savings; serious falls are associated with a financial consequence for the social care sector, and reducing falls is presumed to prevent such costs from being incurred [D]. With strains on public service budgets throughout the country, the county council is interested in alternative strategies to reduce costs.

To satisfy the need to demonstrate the preventative care programme’s financial benefit, the project team dedicated significant effort towards developing an evaluation strategy. Evaluation strategies entail an abundance of choices; each comes with particular trade-offs and subtle considerations. For example, these include determining which measures to collect, what design the evaluation should follow, the number of service users required to determine

whether the programme is effective, and the process and cadence of reviewing the pilot’s status. Here, we focus specifically on the evaluation design and the complicated trade-offs managed by the project team.

Early on, the consultants suggested “an RCT [randomised controlled trial], where people identified by the model are divided into treatment and control groups” [C3]. Yet, following a RCT design presents numerous complications, the most salient of which were associated with the legal and ethical challenges presented by using a control group. The ethical dilemma of using a control group was commonly framed as whether the county council could “not give [service users] a service even if we know they have a need” [CC18]. This question arose at several times throughout the project. First by two consultants who concluded that it was unethical to not intervene on someone who is known to be at risk [C3, C8]; later by a different consultant who intended to use a control group during evaluation; and finally by practitioners planning the implementation of the ML system within the district councils [CC15, CC16, CC18].

The repetition of this debate reflects several complicating features of the field site. First, participants perceived there to be a lack of institutional support that can provide guidance on programme evaluations [CC2, CC5, CC9, C7]. Second, members of the project team consistently changed: consultants “rolled off” to begin working for different clients and practitioners left to take up posts elsewhere in the council. In the words of one stakeholder, “that’s always difficult because you keep having to go over what’s been done and bringing people up to speed” [TP1]. Throughout these transitions, the question of the control group only ever achieved partial closure before being reopened by a new group of stakeholders.

## 4.2 Adapting to Emergent Challenges Through Pragmatic Resolutions

ML system implementations are difficult work. During their efforts, the project team was presented with and adapted to several prominent challenges, shown in Table 3. These obstacles manifest in discussions around data sharing agreements, securing capacity for the services that will be offered through the preventative care programme, conducting rigorous evaluations of a system obfuscated by intellectual property protection, and determining whether the programme is engaging with the “right cohort” of people. In the face of these challenges, stakeholders responded by devising **pragmatic resolutions** – extemporaneously contrived techniques that drew on the resources and expertise at hand to mitigate the effect of these challenges on the implementation process. Pragmatic resolutions may come in conflict with one another, as well as change as stakeholders experiment with different techniques.

**4.2.1 Negotiating Data Sharing Agreements.** An early and critical task in developing the ML system was negotiating data sharing agreements between the different organisations to access the data required to train the ML model. Within any one county, there are many data sources that may be relevant to a project but are distributed across organisational boundaries [C6, TP1]. For example, ASC data may reside with the county council, while revenues and benefits data may be controlled by the district councils. And as per several stakeholders, these local authorities are “very nervous

**Table 3: Four prominent challenges experienced during the implementation of the ML system and their associated pragmatic resolutions. Challenges with multiple pragmatic resolutions indicate instances when stakeholders experimented with different techniques.**

Sec.	Challenge	Description	Pragmatic Resolution(s)
4.2.1	Negotiating Data Sharing Agreements	Disparate data sources and risk averse culture complicate efforts to share data across organisational boundaries.	Data minimisation: Limiting number of data sets and fields consumed by ML system.
4.2.2	Securing Capacity for Prevention Services	Limited capacity among public services to support preventative programme.	Feature importance: Anticipating service allocation through local feature importance. Batching: Contacting service users in sequential batches.
4.2.3	Intellectual Property Impacts System Evaluation	Intellectual property protection limits ability of external organisations to evaluate ML system performance.	Dip checks: Manually comparing input and output data.
4.2.4	Competing Conceptions of the Target Population	Various interpretations of who constitutes the right cohort for the programme.	Exclusion criteria: Excluding ineligible participants after predictions are generated. Data extension: Expanding data sources consumed by ML system.

around sharing data” with one another, or with health organisations and external providers [TP1, CC8]. Stakeholders consistently attributed the reluctance to share data to “mistrust” [D] between the county and district councils and a risk averse culture [TP1, C7, CC2]. Within this context of hesitancy, the activity of negotiating access to data involved efforts to minimise the amount of data shared across organisations: for example, determining the relevant data sources and selecting specific data fields from the CMS to use during model training.

As mentioned, data sets in a given locality are held across numerous organisations depending on the types of services each provides. With prior clients, the technology provider used data from a range of services, such as adult and children’s social care, housing, revenue and benefits, and debt services. Yet, for the pilot phase of this project, the county council decided to only share data from adult social services.

**C8:** “The data sources were limited by what we had from an IG [information governance] perspective. That was the limiting factor. We wanted to use more data from other organisations because we know that would improve performance.”

To access these data, consultants and individuals in the technology provider had to navigate tense data sharing negotiations:

**TP1:** “Every local authority, they’re very nervous around sharing data, so they need a lot of reassurance. People underestimate that, and it takes a lot of time. Once they start learning about what data needs to get shared, the parties involved, and what that all looks like, they get fidgety.”

After agreeing upon the data sets to be shared, the specific fields were then selected. Practitioners in the county council decided to not include protected characteristics in the ML model due to concerns over sharing sensitive information with external organisations. Therefore, service user information such as age, gender, or race are not used as predictors [C2, TP2]. Stakeholders acknowledged the sensitivity of including protected characteristics in the ML model but simultaneously interpreted their exclusion as hindering predictive performance [TP2, C8] and the ability to evaluate the ML model for instances of racial or gender bias [C2].

**4.2.2 Securing Capacity for Prevention Services.** This project’s span across traditional organisational boundaries posed not only difficulties for navigating data sharing agreements, but also for negotiating how the preventative care programme will be delivered. Specifically, a point of contention within the programme, and one that led to a month-long project delay, was deciding what prevention services will be offered by the programme and how many service users each could support. Many of the prevention services the project team hoped to refer service users towards – for example, multifactorial falls assessments conducted by a community health team – are grappling with capacity constraints of their own. Therefore, the project team had to reassure the services that they would not be over-burdened by this newly uncovered demand.

Reassuring the prevention services implied a balance between two competing requirements. At one end, the project team wanted to include enough service users in the programme so they could estimate the effectiveness of each service [C4, C8]. On the other end, the project team did not want to direct too many people to any

one service, overwhelming already strained organisations [CC1, CC2, CC8, C7].

The project team made several attempts to address this problem. Determining capacity required numerous virtual meetings between members of the project team and the leads of each service [CC5, C7]. To support these negotiations, the project team wanted to confirm with each service that they will not receive more referrals than they have capacity to support.

As a first attempt to provide this reassurance, the consultants speculated whether statistical methods could infer which intervention a service user will likely receive prior to the holistic conversation. In collaboration with the data scientist from the technology provider, the group produced local feature importance calculations to indicate the contribution of specific features to each of the ML system's predictions. Initially, the data scientist cautioned against the validity of this strategy, pointing to the objective for which the ML model was developed. Nonetheless, the team experimented with this approach, generating a list of the top fifteen most important features for each prediction. Unfortunately, this strategy failed because the features used in the ML model did not align with the interventions included in the programme: a feature may be important in predicting that a person is at risk of falling, but it does not necessarily correspond to a falls prevention service.

In response to this obstacle, the project team conceived an alternative method. Rather than continuously engaging new service users for the holistic conversations, the population was divided into six batches, each of whom were contacted in intervals. The volume of service users referred to each prevention service was continuously monitored by the project team. If a service nears capacity, the Population Health Contact Centre re-directed service users to alternative services.

**4.2.3 Intellectual Property Impacts System Evaluation.** The technology provider considered the risk framework to be its intellectual property; therefore, understandings of the risk framework among the consultants, county council, and research team were limited. This obfuscation led several consultants to refer to the risk framework as a “black box” [C2, C5, C6] and impacted the project team's ability to conduct rigorous evaluations of the ML system. As a result of this limitation, members of the consultancy devised workaround evaluation methods.

Consultants on the project referred to this workaround method to identify potential errors during the NLP analysis as “*dip checks*.” Their method focused on two potential sources of error: whether the issue arose from the data redaction process – “*pipelines issues*” – or from the NLP component. The following fieldwork vignette [24] illustrates this “*laborious process*” of evaluation conducted by one of the firm's consultants.

*On one monitor, the consultant logged into the client management system; on the other, he accessed the master risk table in the SQL server. One by one, he manually matched the 10-digit service user IDs from the CMS with the pseudonyms generated by the information governance bridge. After “a few painful hours,” the consultant found an issue. Service user IDs were incorrectly matched with one another during the data transfer process.*

*The consultant began to compare the case notes with the master risk table. With a sample of 140 service users, he found that around a third of falls identified by the NLP component could not be identified in the case notes. After noting these two issues, he shared the findings with his team and prepared an email for the technology provider.*

As this vignette shows, intellectual property protection shapes evaluation mechanisms, instigating consultants to devise workarounds to see what is “under the hood.” The lack of visibility due to intellectual property protection and the issues identified led some members of the consultancy to question the quality of the risk framework. Yet, after several rounds of dip checks and communicating the results to the technology provider, the consultants began to notice improvements in the risk framework and consequently adjusted their perceptions. For example, one consultant described how their evaluation of the risk framework changed over time:

*“If you'd asked me five months ago, I'd have said not confident at all...Now there are still a few things missing but we've done a lot of manual checks and we've seen the quality of those checks improve over time.”*

**4.2.4 Competing Conceptions of the Target Population.** A salient concern among the project team was whether the population identified by the ML system was appropriate for the preventative care programme. Several stakeholders conducted various modes of systematic and *ad hoc* evaluations: practitioners cross-referenced samples of predictions with case notes in the CMS, a project manager filtered the predictions by a set of exclusion criteria, and data scientists communicated ML model performance through its precision. Each evaluation surfaced different notions of what constitutes the “right cohort,” leading to various strategies to ensure the preventative care programme engages with the appropriate population.

One definition of “right” was that the service users who are contacted for the holistic conversation are neither deceased nor in residential care, and are over the age of 65. This information was perceived to be readily available in the CMS and could therefore be instantiated as a series of filters in the data pack. Prior to sharing the data pack with the Population Health Contact Centre, a practitioner can filter the predictions according to these three criteria, ensuring that the cohort is “right.”

Meanwhile, the next interpretation suggests that a valid population is composed of those who are not known to adult social services. Practitioners in the county council recognised that people oftentimes do not engage with the social care system until their needs have reached a level of acuity to require formal services. A population with a high level of need was perceived to pose some risk to the programme: either the prevention services would not be able to support them or an excess number of service users would be referred to the already-strained health system. Given that many people known to ASC are already at a high level of need, several practitioners proposed that the “right cohort” may be one that is not known to the system [CC2, CC15, CC18].

**CC2:** *“My biggest concern is that this is the wrong cohort, because these people are already known to us... Do we need to look more towards the district side of the*

*data? If we want to do what the name says – proactive – and reach out to people that aren't known to us."*

At its current design, the ML system can only generate predictions on those who are already known to the county's adult social services department. The reason is that the project team only has access to ASC data; this decision was made early on in the project primarily due to concerns among the information governance team regarding sharing data across organisations and communicating changes in data use to service users [TP1, C6, C8, CC7].

The limited population of service users for whom the system can generate predictions was problematic for some practitioners. Early on, one practitioner reflected on the dissonance between the purpose of the project and the target population: *"I thought the purpose of this project was to prevent people from getting to us. Working with people we already know is a completely different project"* [CC2]. This concern was commonly allayed by the consultants and technology provider through reference to the future direction of the project: data sources beyond ASC will be consumed by the ML system so it can generate predictions on those who have not engaged with the social care system. For example: *"That is going to be part of the next phase, where we begin working with the districts. So, we'll have access to that data"* [C8].

## 5 USING HCI KNOWLEDGE TO ENVISION PREVENTATIVE CARE WORKFLOWS

During our fieldwork, it became apparent that HCI considerations such as identifying a user, understanding their needs, and integrating the system into their workflow were not considered (Section 4.1.3). Therefore, to highlight the value of engaging with HCI-related concerns during the design of ML systems for preventative care programmes, we created an alternative workflow built on HCI knowledge. We base our workflow re-design on two principles and four design guidelines derived from past research, and validate it with stakeholders from the county council, consultancy, and technology provider.

### 5.1 Principles & Guidelines for ML Systems in Preventative Care

Risk and its management have become crucial components of ASC practice over the past three decades [75]. There exist numerous resources for practitioners, many of which emphasise the importance of equal participation between service users and practitioners and accounting for the many personal and environmental circumstances that amplify or mitigate risks [5, 43, 71]. Put simply, risk assessments should be **subjective** and **contextual**.

Apart from this normative position on risk assessments, we draw from HCI and Social Work research to define four guidelines for ML-enabled workflows.

- (1) Design should not lead to the creation of unnecessary **administrative work** that detracts from time spent engaging with service users [12, 36, 50, 84].
- (2) Design should facilitate practitioners' need to construct **narratives** of service users [84, 98].

- (3) Design should facilitate practitioners' need to use **contextual information** when evaluating the validity of predicted risk scores [17, 47, 97].
- (4) Designers should incorporate the system at a point in workflows where practitioners are likely to **naturally encounter** it [106].

### 5.2 Integrating Predicted Risk Scores into Client Management Systems

We propose an alternative workflow depicted in Figure 4. Much is the same as in the current workflow (Figure 2), besides two crucial adjustments in Steps 4 and 5. Rather than preparing a disparate *Excel* spreadsheet at Step 4, the predicted risk score is integrated into the CMS within each service user's profile. This workflow can utilise the customisation afforded by CMSs to include a field for each service user's risk score that is automatically populated each time the predictions are generated; in doing so, the **administrative work** of preparing a separate data pack is obviated. And by residing within service users' profiles, practitioners can review predicted risk scores alongside case notes. This enables practitioners to draw on **contextual information** when assessing the validity of predictions and incorporate the scores into a broader **narrative** for the service user. Finally, as practitioners regularly interact with the CMS, the risk score's inclusion as an automatically generated field should increase opportunities for practitioners to **naturally encounter** it.

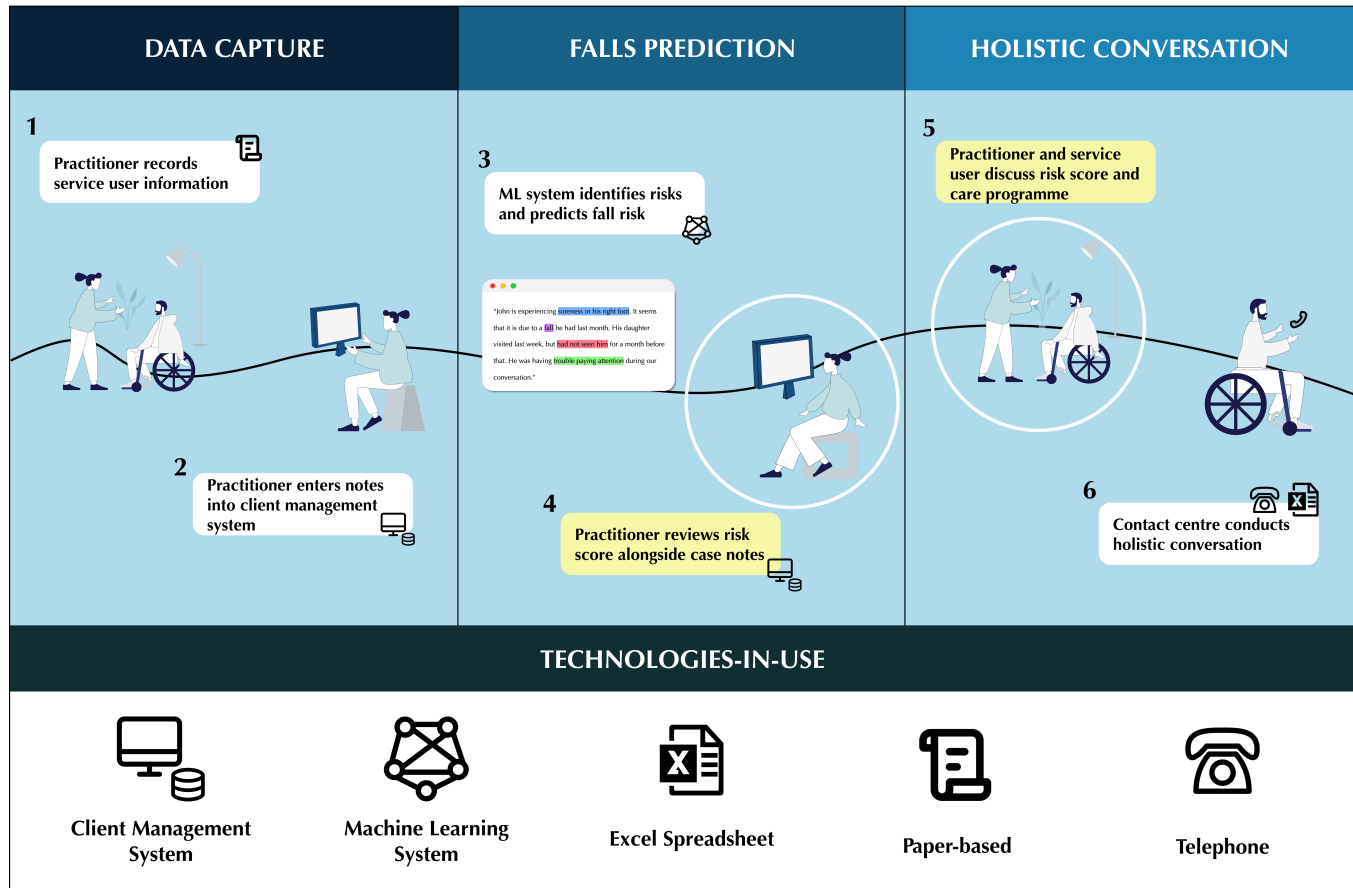
The next change is present in Step 5. Whereas before service users were initially informed of their inclusion in the preventative care programme by a letter sent to their homes, the alternative workflow utilises practitioners to disseminate this information. During subsequent interactions with service users, practitioners discuss the risk of falling – if the risk score sits above a specified threshold – and the preventative care programme. This approach, rather than relying on communication through letters, supports the **subjectivity** and **contextual specificity** of the risk assessment process. First, it provides service users with the opportunity to discuss their risk of falling with a trained practitioner before agreeing to participate in the preventative care programme. Second, as these discussions are held in-person, practitioners can conduct a preliminary assessment of the service user's home environment to contextualise the predicted risk score. While this process does lead to the creation of additional efforts for practitioners, these tasks are, crucially, not administrative: they involve direct interaction with service users.

In summary, the alternative workflow aims to stimulate practitioner-service user conversation on the risk of falling and support interaction opportunities with the predicted risk scores.

### 5.3 Stakeholder Reflection Sessions

To validate the proposed workflow, we conducted stakeholder reflection sessions [38] with individuals from the county council, technology provider, and consultancy. Seven individuals agreed to participate in the sessions: this includes, as shown in Table 4, two from the county council, one from the technology provider, and four from the consultancy. To accommodate stakeholders' schedules and to reduce the demands placed on them, we agreed to conduct the

## ALTERNATIVE WORKFLOW



**Figure 4: The alternative workflow to promote service user engagement. Adjusted steps (4 and 5) are circled and highlighted in yellow. Rather than preparing a separate data pack and contacting service users by letter, predicted risk scores are integrated into the CMS for social care practitioners to review. Practitioners then incorporate discussion on fall risk into routine interactions with service users.**

reflection sessions during the project team’s recurring meetings. In total, two stakeholder reflection sessions were conducted with each lasting approximately 30 minutes.

**Table 4: Overview of stakeholders (n=7) from the county council, technology provider, and consultancy involved in the reflection sessions.**

ID	Organisation	Gender	Session
R1	County Council	F	1
R2	County Council	F	1
R3	Technology Provider	M	1
R4	Consultancy	F	2
R5	Consultancy	F	2
R6	Consultancy	M	2
R7	Consultancy	M	2

The session began with the lead author presenting the principles and guidelines described at the beginning of this section and the alternative workflow in Figure 4. Then, he facilitated a discussion centred on two questions: *what value does this workflow provide?* and *how can the workflow be improved?* The lead author elicited responses from each participant and recorded handwritten notes during these sessions. Notes from the sessions were analysed through affinity diagramming [45]. The presentation shared with stakeholders during this session is available in the Supplementary Material.

## 5.4 Stakeholder Responses

In this section, we present stakeholder responses from the reflection sessions. First, we summarise the value stakeholders perceived in the workflow. Second, we outline the various ways in which the stakeholders perceived it could be improved.

When discussing the value of the workflow in relation to the project, four patterns emerged from stakeholder responses. First,

stakeholders perceived the revised workflow as extensible, supporting the county council's aim of extending into risks other than falls [R1, R2] and enabling practitioners to consider interventions beyond those offered through the holistic conversation [R4]. Second, stakeholders anticipated that front-line practitioners would respond positively to the workflow given that it does not “*prescribe what they should do*” [R3, R4, R7]. Third, members of the consultancy perceived that integrating risk scores with the CMS would positively influence adoption of the ML system [R5, R6, R7]. Finally, stakeholders suggested that the workflow enables the county council to both contact the right cohort [R1] and contact them in the right way [R1, R2, R6]. That is, the workflow allows practitioners to judge, first, whether a service user is appropriate for the preventative care programme during face-to-face interactions and, second, how individual service users would like to discuss the preventative care programme.

Stakeholder discussion on limitations with the revised workflow focused on practical constraints and opportunities for further changes. While valuing the proposal to integrate risk scores in the CMS, two consultants acknowledged the difficulty of doing so within the context of a “*proof-of-concept*” project [R4, R7]. Further, consultants expressed concern that the workflow may lead to additional work for already-burdened social workers [R7] and questioned the assumption within the workflow that social workers have sufficiently frequent interactions with service users to support the programme [R5, R7]. Meanwhile, stakeholders anticipated various opportunities to further improve the workflow. For example, they expressed a desire to support data capture through mechanisms other than handwritten notes [R5, R6], questioned how risk scores could be presented to service users [R5], and proposed the use of explainability techniques to support practitioners' interpretations of risk scores [R4].

## 6 DISCUSSION

### 6.1 Implications for Research

Much research in HCI gravitates towards the two ends of the ML lifecycle. Researchers delineate the design practices of data scientists and their collaborators [e.g., 61, 62, 96, 105, 109]; others instead look towards the experience of use, focusing on instances of interaction with ML systems [e.g., 7, 30, 40, 47, 51, 97]. Conspicuously absent from empirical concern is the messy in-between: the work of implementation. Implementation, in this sense, is understood as the work required to move out of the laboratory, into the field, and into the hands of real-world practitioners.

Through our fieldwork, we highlight the importance of attending to implementation as an empirical site in HCI. In this study, we revealed the strenuous efforts required to put ML systems into practice. Stakeholders from a range of organisations collaborated to conduct the joint work of developing a ML system and instituting organisational change. We also observed how stakeholders adapted in real-time to the challenges that emerge during the implementation process by devising pragmatic resolutions.

In our research, we observed that **implementation required collaboration between a complex network of stakeholders**. Past scholarship has highlighted the interdisciplinary and multi-organisational character of ML projects [e.g., 49, 57, 64, 65, 77].

Yet, these studies tend to focus on collaboration as it relates to the development of a ML system; here, we highlight how such collaboration extends beyond design activities and into the practical work of delivering services enabled by ML systems. Further, the cross-sector participation has implications for long-standing discussions on accountability in computing. Particularly, scholars point to the “problem of many hands”: many people involved in the development, use, and maintenance of software systems challenges efforts to attribute blame [67]. Recently, researchers have highlighted how this problem is exacerbated in ML systems propagated through “algorithmic supply chains” – complex arrangements of actors involved in the production and use of ML systems – and built of innumerable open-source and pre-built modules [18, 100]. Here, the process through which service users are deemed eligible for the preventative care programme is illegible to the county council due to the intellectual property protection guarding the risk framework. Given that the performance of ML systems are sensitive to changes in data distributions [52, 78] and can vary across social groups [4, 68], it becomes essential for social care organisations to have clarity around how ML systems are evaluated and assurance that considerations such as performance degradation and fairness are actively accounted for.

Through our interviews, we found that **these stakeholders were bound together by a multiplicity of motivations. For the county council, the most salient objectives were to realise cost savings by stemming the demand for public services while simultaneously improving service users' welfare**. Yet, as previous work has discussed, evidence for the impact of ML system deployments on a range of outcomes is limited [10, 22, 55]. Further complicating the matter is that research in HCI tends to emphasise the creation of novel artefacts to the neglect of translating those creations into practice [54]. As such, we add to growing calls for further research on the evaluation of ML systems in real world contexts of use [7, 51] by highlighting the challenges associated with ML system implementations.

Additionally, we found that **significant work was required to incorporate the ML system into the preventative care programme**. This unexpected labour took the form of the pragmatic resolutions devised by project stakeholders. This finding confirms the observation presented in numerous other bodies of work that substantial efforts are required to integrate technological systems into working practices [9, 94]. Recently, Elish and Watkins [28] generate the concept of “repair work” to describe the continuous efforts to mend the breakages in workflows and social relationships that accompany the implementation of ML systems. Such accounts complicate narratives of a linear trajectory from technical design to real-world impact, and imply that technologies which are characterised as labour-saving may instead have the opposite effect, at least in the short-term. With the growing concern in HCI to move beyond a preoccupation with “users” and account for the needs of a diversity of stakeholders, future research should explore how HCI can support these critical roles in their efforts to bring ML systems into practice and compensate for the inevitable breakdowns that arise as rigid, technical systems are put to work in the flexible and nuanced world of organisational life [1].

We found a **series of challenges that arose during the course of the project to which stakeholders had to adapt, drawing**

**on the resources and expertise available at the time.** In doing so, we respond to calls for greater clarity around the issues surrounding the adoption of ML systems [58, 105]. With these findings, we contribute to the understanding of challenges associated with ML deployments by both corroborating and extending prior work. First, other studies have found similar issues in the deployment of ML systems, such as navigating regulation and policies during data sharing negotiations [64] and identifying effective interventions to be delivered off the back of predictions [81, 97]. And while past work has pointed to inconsistencies between the population in production and that of model training [e.g., 63], we instead show a tension between the population for whom predictions are generated and the stated goals of the project. Some of the observed challenges remain relevant to the implementation of traditional information systems, signalling that much of what has been learned from past efforts is still applicable to ML systems. Yet, the unique characteristics of ML systems bring particular complexities to prominence. Placed between the data-dependence of ML systems for training and the hesitancy around data sharing across the county, stakeholders were forced to make measured trade-offs on the specific data fields that would be included as predictors in the ML system. The obscuring of the NLP framework by intellectual property protection instigated the consultants to devise a unique evaluation method that explicitly treated the system as a “black box,” comparing inputs to outputs. And concerns over data sharing manifest once again to constrain the perceived utility of the ML system by delimiting the population for whom the system can generate predictions.

Finally, we contribute to research on how such challenges can be overcome by highlighting the role of **pragmatic resolutions**. The pragmatic resolutions put to work on these challenges may at times be in conflict with one another. At early stages in the project, data privacy concerns were assuaged by limiting the scope of data consumed by the ML system; later on, this scope was expanded in hopes that it will lead to improvements in the system’s utility. Further, how practitioners perceive a challenge and attempt to address it are also dynamic. The character of challenges can change over time and with it the pragmatic resolution. When considering what constituted the appropriate target population for the programme, the pragmatic resolution shifted from filtering out ineligible service users to extending the data sources. Since various scholars have argued that interventions to design practices are most effective when they are based on practitioner’s existing techniques [85, 91], we propose that future research should identify how practitioners in a range of organisations attempt to address implementation challenges as a means to inform design interventions.

## 6.2 Guidance for Implementing ML Systems in Social Care

Now, we provide implications for implementing ML systems in the ASC sector in England. We take a broad view, advocating for interventions across varying levels of social and technological organisation.

- **Begin ML projects from a problem-focused orientation.** An undesirable consequence of contemporary interest in AI and ML is that many fall into the trap of technological solutionism, or the belief that these technologies can be used

to address any problem [60]. As a result, some organisations begin with the technology and go off in search of a problem. This issue is particularly acute in the social care context where ML and other digital technologies are repeatedly forwarded as a viable intervention to the many financial and operational challenges afflicting the sector [42, 69, 72]. Instead, organisations should first understand the specific problems they contend with and determine whether ML is an appropriate intervention. The Health AI Partnership provides detailed guidance on the steps practitioners can take to identify and prioritise relevant problems afflicting their organisation [74]. This determination should involve critical questioning of whether it is reasonable to assume that the phenomenon of interest – risk of falling, social isolation, or criminality – can be predicted from historical data [95].

- **Set the foundations for ML projects with effective and efficient data capture mechanisms.** Before moving into the stages of model development, practitioners should determine whether available data contain measures that serve as valid representations of the phenomena of interest [90]. Inadequate proxies – for both features or the target variable – may lead to unintended results. For example, Obermeyer et al. [68] show how a widely-used ML system in US hospitals relied on costs as a proxy for health needs. Yet, because of unequal levels of access to healthcare, black patients tend to have lower costs than their white counterparts. As a result, the ML system assigned lower risk scores to black patients with the same health status as white patients. Therefore, rather than relying on indirect approximations, practitioners should implement data capture mechanisms within existing information systems. Using the current case as an example, an alternative approach to relying on the extraction of “fall risks” from case notes would be to create structured fields for variables such as falls incidents within the CMS.
- **Discuss data sharing early and often.** As found in this study, and noted elsewhere [56, 80], public services – and, as a result, the data each generates – are fragmented across a range of disparate organisations. Integrating these data introduces not only technical challenges, but also entails a commensurate amount of organisational work. As found here, some groups may be reluctant to share data across traditional organisational boundaries. Therefore, it is critical to incorporate relevant stakeholders – from data protection analysts to service managers – into discussions on data sharing as early as possible. Due to the dependence of ML systems on large amounts of data, this need is heightened as compared to traditional information systems implementations.
- **Integrate system and workflow design processes.** A consistent finding in Social Work and HCI research is that users resist technological systems that prescribe workflows which do not align with the contingent and collaborative character of work practices [50, 97, 98, 101]. Aligning ML systems with workflows can follow several directions. One method is to integrate ML systems into existing workflows. Yang et al. [106] demonstrate this strategy by identifying a time and place in clinicians’ workflows at which they are likely to naturally encounter a ML system. Alternatively, the



implementation of ML systems may require the creation of entirely new workflows. Regardless of which approach is taken, decisions on how a workflow should progress has implications that pervade system design considerations, as we illustrate in Section 5.

- **Negotiate system evaluation requirements with third-party providers.** Our findings demonstrate that evaluating third-party systems is complicated by legal frameworks such as intellectual property protection. Further, recent work shows that many social care organisations lack the internal expertise to evaluate ML systems during procurement [53], potentially due to practitioners' lack of experience with the technology [8]. These considerations have a severe implication for the use of ML systems in social care. Namely, the process through which public services are delivered becomes partially obscured to the organisation that has statutory responsibility for ensuring such services are delivered effectively and fairly. Social care organisations should request third-party providers to explain how accuracy of the ML system will be monitored and maintained, including over time and across sub-groups of the target population, how deviations from expected system behaviour will be identified and addressed, and what steps will be taken to comply with regional data protection regulation [79]. Various standards are available which can support these efforts, such as the Algorithmic Transparency Recording Standard [15], Data Protection Impact Assessments [73], and Algorithmic Impact Assessments [46].
- **Establish a robust and ethical approach to programme evaluation.** Our findings show how organisational interests demand a rigorous approach to evaluation to demonstrate whether a programme achieves its aims. Yet, programme evaluation in general is a complex undertaking; this complexity is heightened when focusing on ML-assisted preventative care programmes for several reasons. First, intended outcomes of preventative care take time to manifest. Second, as recognised by stakeholders in this study, use of the ML system uncovered a population of at-risk service users in a manner that was not previously possible. Therefore, adopting an evaluation design such as a randomised controlled trial implies that the county council would temporarily withhold services from those it has identified may be in need of them. These two considerations point to the need for organisations to develop an approach to evaluation that is robust and ethical. Robust implies that the evaluation design can be tailored to suit the circumstances and needs of particular programmes. There are myriad types of evaluations: formative evaluations that seek to improve how programmes are delivered and summative ones that aim to determine a programme's outcomes; randomised controlled trials, comparison group designs, and case studies; or those that collect quantitative and/or qualitative data on a range of economic, operational, and service user outcomes [88, 99]. Meanwhile, an ethical approach to evaluation suggests that organisations recognise and are able to mitigate the ethical implications associated with the various designs.

## 7 CONCLUSION

In this study, we describe the experience of implementing ML systems in social care organisations. Based on longitudinal fieldwork, we reveal the complicated network of stakeholders implicated in catalysing the county council's efforts to transition towards a proactive practice and the diverse range of social, technical, and legal challenges presenting to such projects. We discover the importance of pragmatic resolutions in contending with challenges as they arise during the practical, day-to-day work of system implementations. Through our research, we gain a deeper understanding of implementation challenges for ML systems to support future deployments, recommend changes to the workflows of preventative care programmes, and delineate guidance for implementing ML systems in social care.

## ACKNOWLEDGMENTS

We thank all of the stakeholders from the county council, consultancy, and technology provider for their openness, insight, and patience. We thank Anna Kawakami, Hayoun Noh, and Thomas Serban Von Davier for their critical feedback on earlier drafts of this paper. We thank the anonymous reviewers for their comments which significantly improved this paper. This research is co-funded by a UK-based consultancy and the Oxford-Singapore Human-Machine Collaboration Programme, supported by a gift from Amazon Web Services. The name of the consultancy is redacted to preserve the anonymity of participants involved in this research. The consultancy was not involved in research design, data collection and analysis, results interpretation, or writing. In accordance with the Russell Studentship Agreement, the consultancy retains the right to review papers prior to publication.

## REFERENCES

- [1] Mark S. Ackerman. 2000. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human-Computer Interaction* 15, 2 (2000), 179–203. [https://doi.org/10.1207/S15327051HCI1523\\_5](https://doi.org/10.1207/S15327051HCI1523_5)
- [2] Hugh Alderwick, Charles Tallack, and Toby Watt. 2019. What should be done to fix the crisis in social care? <https://www.health.org.uk/publications/long-reads/what-should-be-done-to-fix-the-crisis-in-social-care>
- [3] Asbjørn Ammitzbøll Flügge, Thomas Hildebrandt, and Naja Holten Møller. 2021. Street-Level Algorithms and AI in Bureaucratic Decision-Making: A Caseworker Perspective. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 40 (2021), 23 pages. <https://doi.org/10.1145/3449114>
- [4] Solon Barocas and Andrew D Selbst. 2016. Big Data's Disparate Impact. *California Law Review* 104, 3 (2016), 671–732. <https://doi.org/10.15779/Z38BG31>
- [5] Monica Barry. 2007. *Effective Approaches to Risk Assessment in Social Work: An International Literature Review*. Technical Report. Scottish Executive Social Research.
- [6] Howard Becker and Blanche Geer. 1957. Participant Observation and Interviewing: A Comparison. *Human Organization* 16, 3 (1957), 28–32.
- [7] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M. Vardoulakis. 2020. A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376718>
- [8] Margaret Blake, Claire Lambert, Rachel Burkitt, David Hills, Freddie Gregory, Laura Tuhou, Stephen Mallet, Jane Stevens, Fiona Richardson, Jonathan Gardam, Paul Riley, Helen Lancaster, Diane Buddery, and Liz Burtney. 2021. *NHSX Adult Social Care Technology and Digital Skills Review*. Technical Report. Ipsos MORI. 230 pages.
- [9] John Bowers. 1994. The Work to Make a Network Work: Studying CSCW in Action. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, North Carolina, USA) (CSCW '94). Association for Computing Machinery, New York, NY, USA, 287–298. <https://doi.org/10.1145/192844.193030>

- [10] Jonathan Bright, Bharath Ganesh, Cathrine Seidelin, and Thomas M. Vogl. 2019. *Data Science for Local Government*. Technical Report. Oxford Internet Institute.
- [11] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-making in Child Welfare Services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, UK). ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300271>
- [12] Judith Burton and Diane Van den Broek. 2009. Accountable and Countable: Information Management Systems and the Bureaucratization of Social Work. *British Journal of Social Work* 39, 7 (2009), 1326–1342.
- [13] Prime Minister's Office Cabinet Office, Department of Health & Social Care. 2021. *Build Back Better: Our Plan for Health and Social Care*. Technical Report. Cabinet Office, Department of Health & Social Care, Prime Minister's Office.
- [14] Oscar Cariceo, Murali Nair, and Jay Lytton. 2018. Data science for social work practice. *Methodological Innovations* 11 (2018), 1–8. <https://doi.org/10.1177/2059799118814392>
- [15] CDDO and CDEI. 2023. Algorithmic Transparency Recording Standard Hub. <https://www.gov.uk/government/collections/algorithmic-transparency-recording-standard-hub>
- [16] Labour & Equalities Centre for International Research on Care. 2018. *Care data: the scope and quality of data on care*. Technical Report. University of Sheffield. <http://circle.group.shef.ac.uk/wp-content/uploads/2018/12/SC-Policy-Perspective-Care-Data.pdf>
- [17] Hao-Fei Cheng, Logan Stapleton, Anna Kawakami, Venkatesh Sivaraman, Yanghui Cheng, Diana Qing, Adam Perer, Kenneth Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. How Child Welfare Workers Reduce Racial Disparities in Algorithmic Decisions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 162, 22 pages. <https://doi.org/10.1145/3491102.3501831>
- [18] Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding Accountability in Algorithmic Supply Chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 1186–1197. <https://doi.org/10.1145/3593013.3594073>
- [19] Comptroller and Auditor General. 2021. *The adult social care market in England*. Technical Report HC 1244. Department of Health and Social Care.
- [20] Consilium Research & Consultancy. 2018. *Scoping study on the emerging use of Artificial Intelligence (AI) and robotics in social care*. Technical Report. Skills for Care.
- [21] Claudia J Coulton, Robert Goerge, and Emily Putnam-Hornstein. 2015. *Harnessing Big Data for Social Good: A Grand Challenge for Social Work*. Technical Report Working Paper No. 11. American Academy of Social Work & Social Welfare. 21 pages.
- [22] Kathrin Cresswell, Margaret Callaghan, Sheraz Khan, Zakariya Sheikh, Hajar Mozaffar, and Aziz Sheikh. 2020. Investigating the use of data-driven artificial intelligence in computerised decision support systems for health and social care: A systematic review. *Health Informatics Journal* 26, 3 (2020), 2138–2147. <https://doi.org/10.1177/1460458219900452>
- [23] John W. Creswell and J. David Creswell. 2018. *Research Design: Qualitative, Quantitative & Mixed Methods Approaches* (5 ed.). SAGE Publications, Inc., Los Angeles, CA.
- [24] Kathleen M DeWalt and Billie R DeWalt. 2011. *Participant Observation: A Guide for Field Workers* (2 ed.). AltaMira Press, Plymouth, UK.
- [25] NHS Digital. 2021. Digital Social Care Pathfinders Programme 2019–21. <https://digital.nhs.uk/services/social-care-programme/digital-social-care-pathfinders-programme-2019-21>
- [26] Emma Dowling. 2021. *The Care Crisis: What Caused It and How Can We End It?* Verso, London and New York.
- [27] Nashwa El-Bendary, Qing Tan, Frédérique C Pivot, and Anthony Lam. 2013. Fall Detection and Prevention for the Elderly: A Review of Trends and Challenges. *International Journal on Smart Sensing and Intelligent Systems* 6, 3 (2013), 1230–1266.
- [28] Madeleine Clare Elish and Elizabeth Anne Watkins. 2022. *Repairing Innovation: A Study of Integrating AI in Clinical Care*. Technical Report. Data & Society.
- [29] NHS England. 2023. AI in social care. <https://transform.england.nhs.uk/ai-lab/explore-all-resources/understand-ai/ai-adult-social-care/>
- [30] Riccardo Fogliato, Alexandra Chouldechova, and Zachary Lipton. 2021. The Impact of Algorithmic Risk Assessments on Human Predictions and Its Analysis via Crowdsourcing Studies. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 428 (2021), 24 pages. <https://doi.org/10.1145/3479572>
- [31] Skills for Care. 2015. 6 reasons why data and information is important when running a successful social care business. <https://www.skillsforcare.org.uk/Documents/NMDS-SC-and-intelligence/Informatics/6-reasons-why-data-is-important.pdf>
- [32] Office for Statistics Regulation. 2020. *Adult Social Care Statistics in England*. Technical Report. [https://osr.statisticsauthority.gov.uk/publication/report-on-](https://osr.statisticsauthority.gov.uk/publication/report-on-adult-social-care-statistics-in-england/)
- [33] The King's Fund. 2023. A short history of social care funding reform in England: 1997 to 2021. <https://www.kingsfund.org.uk/audio-video/short-history-social-care-funding>
- [34] Paul Michael Garrett. 2005. Social work's 'electronic turn': Notes on the deployment of information and communication technologies in social work with children and families. *Critical Social Policy* 25, 4 (2005), 529–553.
- [35] Philip Gillingham. 2019. Decision Support Systems, Social Justice and Algorithmic Accountability in Social Work: A New Challenge. *Practice* 31, 4 (2019), 277–290.
- [36] Philip Gillingham. 2021. Practitioner Perspectives on the Implementation of an Electronic Information System to Enforce Practice Standards in England. *European Journal of Social Work* 24, 5 (2021), 761–771.
- [37] Philip Gillingham and Timothy Graham. 2017. Big Data in Social Welfare: The Development of a Critical Perspective on Social Work's Latest "Electronic Turn". *Australian Social Work* 70, 2 (2017), 135–147. <https://doi.org/10.1080/0312407X.2015.1134606>
- [38] Google. [n.d.]. Stakeholder Review. <https://designsprintkit.withgoogle.com/methodology/phase6-validate/stakeholder-review>
- [39] HM Government. 2014. Care Act 2014. <https://www.legislation.gov.uk/ukpga/2014/23/contents/enacted>
- [40] Ben Green and Yiling Chen. 2019. The Principles and Limits of Algorithm-in-the-Loop Decision Making. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 50 (2019), 24 pages. <https://doi.org/10.1145/3359152>
- [41] Kate Hamblin. 2020. *Care System Sustainability: What role for technology? An evidence review*. Technical Report 3. University of Sheffield.
- [42] Kate Hamblin. 2020. Technology and social care in a digital world: challenges and opportunities in the UK. *Journal of Enabling Technologies* 14, 2 (2020), 115–125.
- [43] John Harris and Vicky White. 2018. *Oxford Dictionary of Social Work & Social Care*. Oxford University Press, Oxford, UK.
- [44] Naja Holten Møller, Irina Shklovski, and Thomas T. Hildebrandt. 2020. Shifting Concepts of Value: Designing Algorithmic Decision-Support Systems for Public Services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) (NordCHI '20). Association for Computing Machinery, New York, NY, USA, Article 70, 12 pages. <https://doi.org/10.1145/3419249.3420149>
- [45] Karen Holtzblatt and Hugh Beyer. 2016. *Contextual Design, Second Edition: Design for Life* (2 ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [46] Ada Lovelace Institute. 2022. *Algorithmic impact assessment: user guide*. Technical Report. <https://www.adalovelaceinstitute.org/resource/aia-user-guide/>
- [47] Anna Kawakami, Venkatesh Sivaraman, Hao-Fei Cheng, Logan Stapleton, Yanghui Cheng, Diana Qing, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. Improving Human-AI Partnerships in Child Welfare: Understanding Worker Practices, Challenges, and Desires for Algorithmic Decision Support. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 52, 18 pages. <https://doi.org/10.1145/3491102.3517439>
- [48] Anna Kawakami, Venkatesh Sivaraman, Logan Stapleton, Hao-Fei Cheng, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. "Why Do I Care What's Similar?" Probing Challenges in AI-Assisted Child Welfare Decision-Making through Worker-AI Interface Design Concepts. In *Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 454–470. <https://doi.org/10.1145/3532106.3533556>
- [49] Sean Kross and Philip Guo. 2021. Orienting, Framing, Bridging, Magic, and Counseling: How Data Scientists Navigate the Outer Loop of Client Collaborations in Industry and Academia. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 311 (2021), 28 pages. <https://doi.org/10.1145/3476052>
- [50] Jenny Lagsten and Annika Andersson. 2018. Use of information systems in social work – challenges and an agenda for future research. *European Journal of Social Work* 21, 6 (2018), 850–862. <https://doi.org/10.1080/13691457.2018.1423554>
- [51] Vivian Lai, Chacha Chen, Q Vera Liao, Alison Smith-Renner, and Chenhao Tan. 2021. Towards a Science of Human-AI Decision Making: A Survey of Empirical Studies. [arXiv:2112.11471](https://arxiv.org/abs/2112.11471)
- [52] Thomas Liao, Rohan Taori, Inioluwa Deborah Raji, and Ludwig Schmidt. 2021. Are We Learning Yet? A Meta Review of Evaluation Failures Across Machine Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. 10 pages. <https://openreview.net/forum?id=mPducS1MsEK>
- [53] Ian Litchfield, Jon Glasby, Sarah Parkinson, Lucy Hocking, Denise Tanner, Bridget Roe, Jennifer Bousfield, et al. 2023. "Trying to Find People to Fit the Tech...": A Qualitative Exploration of the Lessons Learnt Introducing Artificial Intelligence-Based Technology into English Social Care. *Health & Social Care in the Community*, Article 9174873 (2023), 11 pages. <https://doi.org/10.1155/2023/9174873>

- [54] Aaron Lyon, Sean A. Munson, Madhu Reddy, Stephen M. Schueller, Elena Agapie, Svetlana Yarosh, Alex Dopp, Ulrica von Thiele Schwarz, Gavin Doherty, Andrea K Graham, Kaylee Payne Kruzan, and Rachel Kornfield. 2023. Bridging HCI and Implementation Science for Innovation Adoption and Public Health Impact. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). Association for Computing Machinery, New York, NY, USA, Article 336, 7 pages. <https://doi.org/10.1145/3544549.3574132>
- [55] David Maguire, Matthew Honeyman, Deborah Fenney, and Joni Jabbal. 2021. *Shaping the future of digital technology in health and social care*. Technical Report. The King's Fund.
- [56] Fola Malomo and Vania Sena. 2017. Data Intelligence for Local Government? Assessing the Benefits and Barriers to Use of Big Data in the Public Sector: Data Intelligence for Local Government. *Policy & Internet* 9, 1 (2017), 7–27. <https://doi.org/10.1002/poi3.141>
- [57] Yaoli Mao, Dakuo Wang, Michael Muller, Kush R. Varshney, Ioana Baldini, Casey Dugan, and Aleksandra Mojsilović. 2019. How Data Scientists Work Together With Domain Experts in Scientific Collaborations: To Find The Right Answer Or To Ask The Right Question? *Proc. ACM Hum.-Comput. Interact.* 3, GROUP, Article 237 (2019), 23 pages. <https://doi.org/10.1145/3361118>
- [58] Sean McGregor. 2021. Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. Association for the Advancement of Artificial Intelligence, Washington DC, USA, 15458–15463. <https://doi.org/10.1609/aaai.v35i17.17817>
- [59] Matthew B. Miles, A. Micheal Huberman, and Saladaña. 2019. *Qualitative Data Analysis: A Methods Sourcebook* (4 ed.). SAGE Publications, Los Angeles, CA, USA.
- [60] Evgeny Morozov. 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. Public Affairs, New York.
- [61] Michael Muller, Ingrid Lange, Dakuo Wang, David Piorkowski, Jason Tsay, Q. Vera Liao, Casey Dugan, and Thomas Erickson. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, UK). Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3290605.3300356>
- [62] Michael Muller, Christine T. Wolf, Josh Andres, Michael Desmond, Narendra Nath Joshi, Zahra Ashktorab, Aabhas Sharma, Kristina Brimijoin, Qian Pan, Evelyn Duesterwald, and Casey Dugan. 2021. Designing Ground Truth and the Social Life of Labels. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 94, 16 pages. <https://doi.org/10.1145/3411764.3445402>
- [63] Dennis H Murphree, Patrick M Wilson, Shusaku W Asai, Daniel J Quest, Yaxiong Lin, Piyush Mukherjee, Nirmal Chhugani, Jacob J Strand, Gabriel Demuth, David Mead, et al. 2021. Improving the delivery of palliative care through predictive modeling and healthcare informatics. *Journal of the American Medical Informatics Association* 28, 6 (2021), 1065–1073.
- [64] Nadia Nahar, Haoran Zhang, Grace Lewis, Shurui Zhou, and Christian Kästner. 2023. A Meta-Summary of Challenges in Building Products with ML Components – Collecting Experiences from 4758+ Practitioners. , 15 pages. arXiv:2304.00078
- [65] Nadia Nahar, Shurui Zhou, Grace Lewis, and Christian Kästner. 2022. Collaboration challenges in building ML-enabled systems: communication, documentation, engineering, and process. In *Proceedings of the 44th International Conference on Software Engineering* (Pittsburgh, Pennsylvania) (ICSE '22). Association for Computing Machinery, New York, NY, USA, 413–425. <https://doi.org/10.1145/3510003.3510209>
- [66] NICE. 2019. Falls in older people: assessing risk and prevention. <https://www.nice.org.uk/guidance/cg161>
- [67] Helen Nissenbaum. 1994. Computing and accountability. *Commun. ACM* 37, 1 (1994), 72–80. <https://doi.org/10.1145/175222.175228>
- [68] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [69] Department of Health & Social Care. 2022. Data saves lives: reshaping health and social care with data. <https://www.gov.uk/government/publications/data-saves-lives-reshaping-health-and-social-care-with-data/data-saves-lives-reshaping-health-and-social-care-with-data>
- [70] Department of Health and Social Care. 2021. *People at the Heart of Care: Adult Social Care Reform White Paper*. Technical Report.
- [71] Royal College of Occupational Therapists. 2018. *Embracing risk; enabling choice: Guidance for occupational therapists*. Technical Report. Royal College of Occupational Therapists.
- [72] Institute of Public Care. 2021. *Adoption and Scalability of Technology Innovation in the Adult Social Care Sector*. Technical Report. NHSX. <https://ipc.brookes.ac.uk/publications/building-the-evidence-base-for-tech-innovation-in-adult-social-care>
- [73] Information Commissioner's Office. 2018. Data protection impact assessments. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/guide-to-accountability-and-governance/accountability-and-governance/data-protection-impact-assessments/>
- [74] Health AI Partnership. 2023. Key Decision Points: 8 Decision Points to Consider When Implementing an AI Solution. <https://healthaipartnership.org/key-decisions-in-adopting-an-ai-solution>
- [75] Nigel Parton. 1996. Social work, risk and 'the blaming system'. In *Social Theory, Social Change and Social Work*. Routledge, London, UK, 98–114.
- [76] Nigel Parton. 2008. Changes in the Form of Knowledge in Social Work: From the 'Social' to the 'Informational'? *British Journal of Social Work* 38, 2 (2008), 253–269.
- [77] David Piorkowski, Soya Park, April Yi Wang, Dakuo Wang, Michael Muller, and Felix Portnoy. 2021. How AI Developers Overcome Communication Challenges in a Multidisciplinary Team: A Case Study. *Proceedings of the ACM on Human-Computer Interaction* 5 (2021), 1–25. Issue CSCW1. <https://doi.org/10.1145/3449205>
- [78] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 959–972. <https://doi.org/10.1145/3531146.3533158>
- [79] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 33–44.
- [80] Helene Friis Ratner and Kasper Elmholt. 2023. Algorithmic constructions of risk: Anticipating uncertain futures in child protection services. *Big Data & Society* 10, 2 (2023), 12 pages.
- [81] Santiago Romero-Brufau, Kirk D Wyatt, Patricia Boyum, Mindy Mickelson, Matthew Moore, and Cheristi Cognetta-Rieke. 2020. Implementation of Artificial Intelligence-Based Clinical Decision Support to Reduce Hospital Readmissions at a Regional Hospital. *Applied Clinical Informatics* 11, 4 (2020), 570–577.
- [82] Devansh Saxena, Karla Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 348 (2021), 41 pages. <https://doi.org/10.1145/3476089>
- [83] SCIE. 2020. *Beyond COVID: New thinking on the future of adult social care*. Technical Report. <https://www.scie.org.uk/care-providers/coronavirus-covid-19/beyond/adult-social-care>
- [84] I. Shaw, M. Bell, I. Sinclair, P. Sloper, W. Mitchell, P. Dyson, J. Clayden, and J. Rafferty. 2009. An Exemplary Scheme? An Evaluation of the Integrated Children's System. *British Journal of Social Work* 39, 4 (2009), 613–626. <https://doi.org/10.1093/bjsw/bcp040>
- [85] Mona Sloane and Janina Zakrzewski. 2022. German AI Start-Ups and "AI Ethics": Using A Social Practice Lens for Assessing and Implementing Socio-Technical Innovation. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 935–947. <https://doi.org/10.1145/3531146.3533156>
- [86] Logan Stapleton, Min Hun Lee, Diana Qing, Marya Wright, Alexandra Chouldechova, Ken Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders. In *FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 1162–1177. <https://doi.org/10.1145/3531146.3533177>
- [87] Ruth Thorlby, Ann Starling, Catherine Broadbent, and Toby Watt. 2018. *What's the problem with social care, and why do we need to do better?* Technical Report. The Health Foundation, the Institute for Fiscal Studies, The King's Fund, and the Nuffield Trust.
- [88] HM Treasury. 2020. *Magenta Book: Central Government guidance on evaluation*. [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/879438/HMT\\_Magenta\\_Book.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/879438/HMT_Magenta_Book.pdf)
- [89] Carers UK. 2021. *State of Caring 2021: A snapshot of unpaid care in the UK*. Technical Report. Carers UK.
- [90] Niels van Berkel, Maura Bellio, Mikael B. Skov, and Ann Blandford. 2023. Measurements, Algorithms, and Presentations of Reality: Framing Interactions with AI-Enabled Decision Support. *ACM Trans. Comput.-Hum. Interact.* 30, 2, Article 32 (mar 2023), 33 pages. <https://doi.org/10.1145/3571815>
- [91] Ibo van de Poel, Lotte Asveld, Steven Flipse, Pim Klaassen, Zenlin Kwee, Maria Maia, Elvio Mantovani, Christopher Nathan, Andrea Porcari, and Emad Yaghmaei. 2020. Learning to do responsible innovation in industry: six lessons. *Journal of Responsible Innovation* 7, 3 (2020), 697–707.

- [92] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3174014>
- [93] Johanna Viitanen, Mari Tyllinen, Elina Tynkkynen, and Tinja Lääveri. 2022. Usability of information systems: Experiences of outpatient physicians, outpatient nurses, and open care social welfare professionals from three large cross-sectional surveys in Finland. *International Journal of Medical Informatics* 165 (2022), 104836.
- [94] Lauren Waardenburg, Anastasia Sergeeva, and Marleen Huysman. 2018. Hotspots and Blind Spots: A Case of Predictive Policing in Practice. In *Living with Monsters? Social Implications of Algorithmic Phenomena, Hybrid Agency, and the Performativity of Technology*. Springer Cham, 96–109.
- [95] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2023. Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms that Optimize Predictive Accuracy. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 626. <https://doi.org/10.1145/3593013.3594030>
- [96] April Yi Wang, Anant Mittal, Christopher Brooks, and Steve Oney. 2019. How Data Scientists Use Computational Notebooks for Real-Time Collaboration. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 39 (nov 2019), 30 pages. <https://doi.org/10.1145/3359141>
- [97] Dakuo Wang, Liuping Wang, Zhan Zhang, Ding Wang, Haiyi Zhu, Yvonne Gao, Xiangmin Fan, and Feng Tian. 2021. “Brilliant AI Doctor” in Rural Clinics: Challenges in AI-Powered Clinical Decision Support System Deployment. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 697, 18 pages. <https://doi.org/10.1145/3411764.3445432>
- [98] David Wastell and Sue White. 2014. Making sense of complex electronic records: Socio-technical design in social care. *Applied Ergonomics* 45, 2 (2014), 143–149. <https://doi.org/10.1016/j.apergo.2013.02.002>
- [99] Joseph S. Wholey, Harry P. Hatry, and Kathryn E. Newcomer. 2010. *Handbook of Practical Program Evaluation* (3 ed.). Jossey-Bass, San Francisco, CA.
- [100] David Gray Widder and Dawn Nafus. 2023. Dislocated Accountabilities in the “AI Supply Chain”: Modularity and Developers’ Notions of Responsibility. *Big Data & Society* 10, 1 (2023), 12 pages. <https://doi.org/10.1177/20539517231177620>
- [101] Christine Wolf and Jeanette Blomberg. 2019. Evaluating the Promise of Human-Algorithm Collaborations in Everyday Work Practices. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 143 (nov 2019), 23 pages. <https://doi.org/10.1145/3359245>
- [102] James Wright. 2020. *Technology of social care: Review of the UK policy landscape*. Technical Report 2. Centre for International Research on Care, Labour & Equalities. 18 pages.
- [103] James Wright and Kate Hamblin. 2023. Technology and adult social care in England. In *Care Technologies for Ageing Societies: An International Comparison*. Policy Press, Bristol, UK, 18–48.
- [104] Marvin N. Wright, Sasmita Kusumastuti, Laust H. Mortensen, Rudi G. J. Westendorp, and Thomas A. Gerds. 2021. Personalised need of care in an ageing society: The making of a prediction tool based on register data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 184, 4 (2021), 1199–1219. <https://doi.org/10.1111/rssa.12644>
- [105] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-Examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376301>
- [106] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300468>
- [107] Qian Yang, John Zimmerman, Aaron Steinfeld, Lisa Carey, and James F. Antaki. 2016. Investigating the Heart Pump Implant Decision Process: Opportunities for Decision Support Tools to Help. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 4477–4488. <https://doi.org/10.1145/2858036.2858373>
- [108] Katri Ylönen. 2023. The use of Electronic Information Systems in social work. A scoping review of the empirical articles published between 2000 and 2019. *European Journal of Social Work* 26, 3 (2023), 1–14. <https://doi.org/10.1080/13691457.2022.2064433>
- [109] Amy X. Zhang, Michael Muller, and Dakuo Wang. 2020. How Do Data Science Workers Collaborate? Roles, Workflows, and Tools. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 22 (2020), 23 pages. <https://doi.org/10.1145/3392826>