

Why Do Decision Makers (Not) Use AI? A Cross-Domain Analysis of Factors Impacting AI Adoption

Rebecca Yu¹, Valerie Chen¹, Ameet Talwalkar¹, Hoda Heidari¹

¹Carnegie Mellon University, Machine Learning Department
{rebyu, valeriechen, talwalkar, hheidari}@cmu.edu

Abstract

Growing excitement around deploying AI across various domains calls for a careful assessment of how human decision-makers interact with AI-powered systems. In particular, it is essential to understand when decision-makers voluntarily choose to consult AI tools, which we term *decision-maker adoption*. We interviewed experts across four domains—medicine, law, journalism, and the public sector—to explore current AI use cases and perceptions of adoption. From these interviews, we identify key factors that shape decision-maker adoption of AI tools: the decision-maker’s background, perceptions of the AI, consequences for the decision-maker, and perceived implications for other stakeholders. We translate these factors into an AI adoption sheet to analyze how decision-makers approach adoption choices through comparative, cross-domain case studies, highlighting how our factors help explain inter-domain differences in adoption. Our findings offer practical guidance for supporting the responsible and context-aware deployment of AI by better accounting for the decision-maker’s perspective.

1 Introduction

Increased interest in deploying practical AI tools for decision-makers has spurred growing research into how human decision-makers interact with these tools, particularly regarding decision-maker reliance on AI (Lee et al. 2020; Bansal et al. 2021; Chen et al. 2023)—that is, the rate at which a human decision-maker replaces their initial judgments with AI recommendations. Such studies typically involve controlled evaluations of AI prototypes in lab settings (Cai et al. 2019; Silva et al. 2024) or pilot deployments of AI tools in real-world environments (Jo et al. 2023; Kawakami et al. 2022), where decision-makers are presented with an AI tool and researchers measure the extent to which they defer to AI guidance. However, AI reliance presupposes decision-maker adoption; in practice, decision-makers may opt not to use these tools at all (*e.g.*, a doctor may not use a model without peer endorsement (Henry et al. 2022)).

As such, we emphasize the importance of studying *decision-maker adoption* of AI—the voluntary and consistent use of an AI tool within a decision-maker’s workflow.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

AI Adoption Sheet

- **Decision-maker background.**
 - How does a decision-maker’s professional experience inform their adoption behaviors?
 - Does the decision-maker have any personal biases that may impact their adoption choices?
- **Perceptions of the AI model.**
 - What flaws does the decision-maker perceive the model they are deciding to consult to have?
 - Do they have any doubts about model trustworthiness or capabilities?
 - What do they believe the capabilities of the model and the benefits of AI adoption are?
- **Consequences for the decision-maker.**
 - What liabilities may the decision-maker face (*e.g.*, legal, professional, political)?
 - Are there uncertainties regarding consequences?
 - How does the decision-maker’s workload change?
- **Perceived implications for other stakeholders.**
 - How do decision-makers believe organizations are impacted by AI adoption?
 - How do decision-makers believe decision-subjects are impacted by AI adoption?

Figure 1: We introduce an AI adoption sheet to highlight factors influencing decision-maker adoption.

Prior work on AI adoption has primarily focused on the perspective of *organizational adoption choices* (Sun and Medaglia 2019; Chatterjee et al. 2021; Wu et al. 2024; Neumann, Guirguis, and Steiner 2024), namely, why an organization may choose to adopt an AI product or service. For example, such work may involve surveys to investigate why a hospital might choose to procure and deploy an AI tool (Poon et al. 2025), potentially irrespective of whether it is used by clinicians down the line. We build on this line of research, but shift the focus to the perspective of *decision-makers* within organizations that have already made AI available to their members. While a few prior stud-

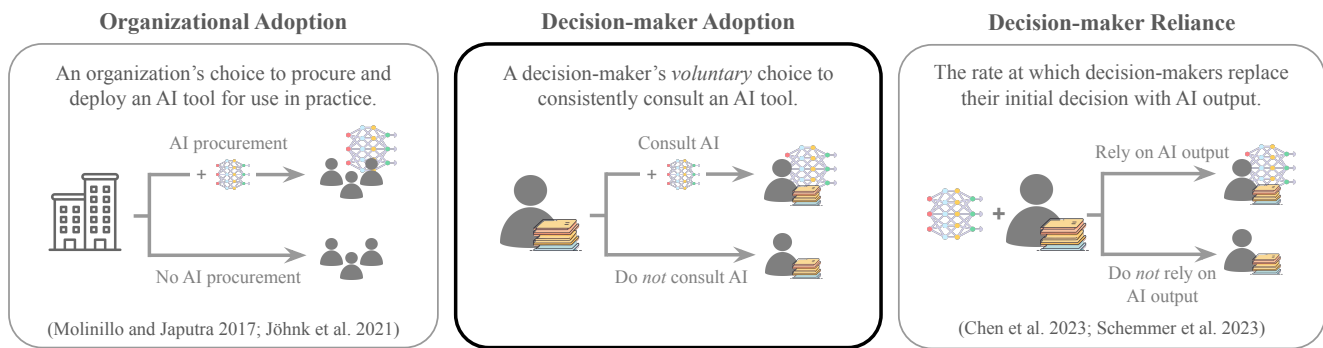


Figure 2: Key definitions to distinguish between organizational adoption, decision-maker adoption, and decision-maker reliance. Adoption refers to an organization or decision-maker’s *choice* to procure or consult an AI tool. In contrast, decision-maker reliance is the rate at which decision-makers replace their initial decision with AI output, which presupposes adoption.

ies have taken a more decision-maker-centric approach in specific domains—primarily medicine (Hameed et al. 2023; Fang et al. 2023; Başar and Erkul 2024), our goal is to identify cross-cutting factors that affect such decision-makers’ adoption of AI tools.

In this work, we conduct a cross-domain analysis to identify inter-domain trends of decision-maker AI adoption choices. We select four representative domains (journalism (Verma 2024), law (Contini 2024), medicine (Park et al. 2019), and the public sector (Amiri and Karahanna 2022)) to cover a variety of predictive and generative AI applications across different sectors of the economy, with varying stakes, and include decision-makers from diverse backgrounds. In total, we conducted 16 semi-structured interviews with decision-makers on their familiarity with AI tools and their considerations for adoption.

Through our interviews, we identify four common key factors that impact decision-maker adoption across domains: (1) the decision-maker’s background, (2) their perceptions of the AI model, (3) consequences for the decision-maker, and (4) the perceived implication of AI adoption for other stakeholders. Through the lens of these four factors, we examine why decision-maker adoption choices for similar AI use cases vary across domains. In particular, we provide two case studies of AI use cases surfaced by participants. First, we discuss why AI is often adopted as an *e-discovery* tool to help gather, organize, and analyze relevant digital data in journalism, law, and the public sector, but not necessarily in medical contexts. Then, we explore why decision-makers in law choose not to use AI for *tailoring communications*, while decision-makers in other domains do. These case studies highlight inter-domain differences in adoption behavior and demonstrate how our proposed four-factor framework can explain when decision-makers will (not) adopt AI.

We envision our framework to serve as a guide to facilitate the proactive assessment of AI proposals for uptake in practice. To provide a structured guide to help researchers, model developers, and model designers think more critically about the key adoption considerations of decision-makers, we turn our factors into an *AI adoption sheet* to help facilitate the integration of decision-maker perspectives into

the model development process (Figure 1). This sheet can serve as a starting point for organizations and policymakers to develop a comprehensive understanding of how decision-makers’ AI adoption choices may manifest in practice. We leave the empirical assessment of the efficacy and usability of the proposed sheet as a critical direction for future work.

2 Related Work

As AI-powered systems are increasingly considered for real-world workflows, research on how decision-makers interact with AI has grown rapidly. We begin by reviewing prior work on *decision-maker reliance*, as it is the most frequently studied aspect of these interactions. However, for decision-makers to rely on AI, they must first choose to consult an AI tool regularly—which we refer to as their *adoption choice*. As such, we then discuss work on AI adoption, highlighting the difference between organizational and decision-maker adoption choices (Figure 2).

2.1 Decision-maker Reliance on AI

A key goal of AI-assisted decision-making is to foster *appropriate reliance*: ensuring that decision-makers can discern when to rely on AI recommendations and when to rely on their own expertise (Schemmer et al. 2023). In practice, simply optimizing for algorithmic accuracy does not guarantee appropriate reliance choices by decision-makers (Cai et al. 2019). Empirical studies have documented patterns of both underreliance (also termed algorithmic aversion) (Dietvorst, Simmons, and Massey 2015; Logg, Minson, and Moore 2019; Cai et al. 2019) and overreliance, where users accept flawed AI outputs (Bansal et al. 2019; Passi and Vorvoreanu 2022). These patterns have motivated work exploring how and when AI outputs influence human decisions (Reverberi et al. 2022; Schelble et al. 2024; Schoeffer, De-Arteaga, and Kuehl 2024).

To mitigate misaligned reliance choices, researchers have proposed interventions such as providing more comprehensive information on model capabilities and limitations (Cai et al. 2019; Yang, Steinfeld, and Zimmerman 2019; Thieme et al. 2023). More recently, studies have moved beyond controlled settings to examine decision-maker reliance choices

Domain	Alias	Domain Experience (Years)	AI Familiarity	AI Interaction Frequency	Primary Area of Expertise
JRN	P01	4	Expert	Daily	Research
	P02	20+	Beginner	Occasionally	Domain
	P03	20+	Intermediate	Frequently	Domain
	P04	10-19	Advanced	Occasionally	Domain
LEG	P05	20+	Intermediate	Frequently	Domain
	P06	10-19	Advanced	Daily	Research
MED	P07	20+	Intermediate	Daily	Domain
	P08	10-19	Advanced	Frequently	Domain
	P09	5-9	Advanced	Daily	Domain
	P10	10-19	Expert	Daily	Research
	P11	10-19	Intermediate	Occasionally	Domain
	P12	20+	Intermediate	Occasionally	Domain
PHW	P13	10-19	Expert	Daily	Research
	P14	10-19	Expert	Daily	Research
	P15	10-19	Beginner	Daily	Domain
	P16	20+	Intermediate	Occasionally	Domain

Table 1: Summary of interview participants along with their relevant domain experience and comfort level with AI (Beginner = Some basic understanding, but limited knowledge, Intermediate = Solid foundational knowledge, Advanced = Deep understanding of core concepts, Expert= Extensive knowledge and experience).

in real-world deployments of AI tools (Kawakami et al. 2022; Jo et al. 2023). Further, as researchers study decision-maker reliance choices in real-world settings, they have highlighted an adoption gap, where the actual adoption and use of models lag behind their development (Dwivedi et al. 2021; Sethi et al. 2020; Uren and Edwards 2023; Cubric 2020). Our work aims to uncover the underlying causes of such lags from the perspective of decision-makers.

2.2 Adoption of AI in Practice

Organizational adoption. Prior work on AI adoption has primarily focused on *organizational* adoption, examining how factors such as technology readiness and innovation culture shape adoption choices (Harvey and Gowda 2021; Uren and Edwards 2023; Campion et al. 2022). Through a combination of surveys, interview studies, and literature reviews on AI deployment, researchers have explored the drivers and challenges of organizational adoption in various domains—including medicine (Davenport and Glaser 2022), business and management (Cubric 2020), construction (Regona et al. 2022), and public administration (Madan and Ashok 2023)). However, findings across these studies can be fragmented, as each study typically focuses on a single distinct domain. We include a brief survey of these factors in the Appendix.

Across ten representative papers (Cubric 2020; Harvey and Gowda 2021; Regona et al. 2022; Bedué and Fritzsche 2022; Campion et al. 2022; Madan and Ashok 2023; Merhi 2023; Uren and Edwards 2023; Neumann, Guirguis, and Steiner 2024; Polisetty et al. 2024), recurring factors include: issues with data quality or availability (10/10); lack

of appropriate data governance and regulation (5/10); and organizational goals to optimize cost or time savings (4/10). While these organizational factors shape the deployment of AI tools and understanding such organizational adoption choices is important, the final choice to use AI tools rests with the individual *decision-makers* responsible for integrating these tools into daily workflows. These decision-makers may choose *not* to use AI, even in the face of organizational support or pressure (Kawakami et al. 2024). We will revisit these factors in Section 6.2.

Decision-maker adoption. Far fewer studies directly address adoption from the perspective of the decision-maker, and those that do often focus on a single domain or use case. For example, Hameed et al. (2023) determined the intentions of healthcare professionals to adopt new technologies through a cross-sectional survey. They predefined seven constructs, which they evaluated to either directly or indirectly influence behavioral intentions. These factors are: performance expectancy, effort expectancy, personal innovativeness in information technology, task complexity, technological characteristics, initial trust, and social influence. In legal use cases, Fang et al. (2023) leveraged semi-structured interviews to understand the attitudes of legal professionals towards AI and identify factors that influence how they engage with AI tools. Their findings revealed that legal professionals and semi-professionals often hold contradictory attitudes towards AI tools. We build on and extend this prior work by looking at AI adoption across multiple domains from the perspective of decision-makers. By identifying cross-domain factors, we aim to draw broader conclusions that can inform future AI deployments in new domains. In Sec-

tion 6.2, we contrast our findings with factors identified in prior work to show why decision-maker perspectives and a cross-domain analysis are essential for understanding real-world AI adoption choices.

3 Study Design and Methodology

To understand what factors influence decision-makers' adoption choices, we conducted an interview study with 16 individuals across four professional domains. Below, we describe our participant recruitment process, study procedures, and the analysis approaches used to synthesize interview data. Participation in the study was voluntary, and the protocol was IRB-approved. We provide further details on our protocol in the Appendix.

3.1 Participants

We focus on four domains: journalism (JRN), the legal profession (LEG), medicine (MED), and the public sector (PHW)—encompassing both public health and welfare. These domains were selected based on three criteria. First, each domain has prior research on AI applications and adoption, enabling participants to speak about concrete, real-world use (or disuse) of AI tools. Second, the domains together must represent a broad spectrum of AI applications, spanning both traditional predictive models and newer generative systems. Third, domains must include participants working in varied professional contexts to capture a breadth of decision-makers. We group “public health” and “welfare” into a composite domain (PHW), as both operate on population-level interventions and share similar methodological approaches.

From each domain, we recruited individuals in the United States who fell into one of two categories: (1) practitioners with professional experience in one of the domains and direct engagement with AI tools, and (2) researcher-insiders who have worked closely and extensively with such practitioners. Including the second category of participants helped mitigate recruitment challenges by broadening the pool of eligible participants while ensuring that all participants could speak meaningfully to the real-world use (or disuse) of AI by decision-makers.

To recruit participants, we began by compiling a list of potential candidates, identified through collaborators in each domain and by locating individuals who had published articles—whether research or opinion—about their experience with AI in their respective domains. We then applied snowball sampling (Goodman 1961; Parker, Scott, and Geddes 2019) to expand our reach and identify additional candidates. We contacted a minimum of 10 individuals within each domain. Despite our effort to diversify interviews across domains, we found that response rates varied substantially: we received 4 responses out of 13 contacts from PHW, 6 out of 24 for MED, 2 out of 11 for LEG, and 4 out of 10 from JRN.

3.2 Semi-structured Interviews

We conducted 16 semi-structured interviews led either by the first and second author (n=7) or one-on-one with the first

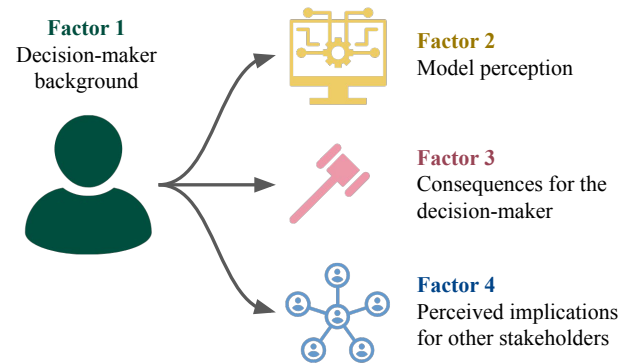


Figure 3: We identified four factors that influence AI adoption choices across domains. All factors are centered around the decision-maker and shape how and why decision-makers choose to adopt AI tools in practice.

author (n=9). The interviews were scheduled via email, and participants completed a consent form before the interview. All interviews were conducted over Zoom and lasted between 45-60 minutes. Participants received a \$100 gift card as compensation.

At the start of each interview, we asked participants to describe their current role and domain experience. We also explored their level of familiarity with AI and the extent of their interaction with AI tools within their domain. Participants were then invited to discuss AI use cases they had either encountered or were otherwise familiar with. Questions were intentionally left open-ended, encouraging discussion of a range of AI use cases—simple or complex. For each use case discussed, we asked follow-up questions to better understand how AI tools were applied in practice, how decision-makers interacted with them, and what considerations influenced decision-makers to choose (not) to use AI tools. Additional details about participant backgrounds can be found in Table 1.

3.3 Analysis Approach

We recorded and transcribed all interviews via Zoom in addition to taking observational notes during each session. Following each interview, we reviewed and edited audio transcriptions for accuracy. We conducted open-coding of the interview transcripts and observational notes, followed by a reflexive thematic analysis (Braun and Clarke 2006, 2019, 2021). Through multiple iterations, we developed and refined the set of factors presented in Section 4. In parallel, we created codes to capture AI use cases mentioned by participants, detailed in the Appendix, which informed our discussion of prevalent use cases in Section 5). Quotations in this paper have been lightly edited for clarity and readability.

4 Factors Impacting AI Adoption

In our cross-domain analysis, we identified four key factors that shape why decision-makers may or may not choose to adopt AI tools (Figure 3).

4.1 Decision-maker Background

The first factor focuses on the decision-makers themselves. In particular, a decision-maker's **professional experience** influences how they engage with AI support tools. Across our interviews, we found that decision-makers often prioritize their professional expertise in AI-generated recommendations, especially in contexts requiring nuanced human judgment. Several participants echoed *"the best clinicians personalize their decisions around the patient"* rather than focusing on what may objectively be the best based on various statistics (e.g., AI recommendations) (P08). One reason is that years of professional experience can provide additional context not captured by models but critical to decision or task outcomes. As P14 noted in PHW, recommendations from AI tools are *"just another piece of information for decision-makers to consider"*, whereas *"people are usually more swayed by personal stories, their own interpretations, and their past experiences"* (P14).

Beyond professional experience, decision-makers' **personal biases** may also play a role in adoption choices. For example, some decision-makers may prefer tools that support rather than replace them. P16 mentioned that while they occasionally use AI tools for brainstorming, *"I have a little bit of bias—or perhaps it's how I was schooled—but I do believe in authoring everything I write myself."*

4.2 Model Perception

A decision-maker's perception of AI tools also impacts their adoption choices. Participants' concerns clustered around three main areas: concerns about model flaws, perception of model transparency and trustworthiness, and faulty beliefs in model capabilities.

Concerns about model flaws. Many participants noted that AI tools often underperform and fall short of expectations. For example, in a PHW use case, P15 described AI-generated social media content as *"comparable to an intern level and something to work with and play with"*, but not mature enough that decision-makers *"would feel comfortable putting out there just yet."* Concerns about *"unexplained phenomena like hallucinations"* or opaque black-box mechanisms were widespread amongst participants (P02). Hallucinations are particularly concerning in public health. P15 specifically mentioned *"We have to be very careful about what we say and how we say it to not feed misinformation, perpetuate it, or misguide people"*. Some participants described the improper use of AI as *"disturbing"* (P07) and alarming to decision-makers (P05), noting that decision-makers have become highly cautious to the point that *"it is now a little shameful to even just use ChatGPT for legal work"* (P06).

However, perceptions can change over time. Participants in all our domains have described increasingly positive perceptions of AI tools with increased exposure to models, especially if there is also organizational adoption and support:

"I think that initially, clinicians were not into AI tools, but I think with the new technologies and interfaces that now exist, I'm starting to see clinicians trying it out and being more interested. [Institution] recently

launched a secure form of a language model hospital-wide for clinicians to try, and I know people are trying that out." (P09)

Perception of model transparency and trustworthiness.

Several participants mentioned that model decisions are significantly more opaque than human decisions, with the lack of explanations making mistakes harder to identify (P01, P05, P06, P08, P09, P14, P15). Improper decisions or outputs from AI tools—especially when uncaught—*"can turn out to be very costly. Humans make mistakes too, but when a human makes a mistake, you have various justifications for it, such as someone didn't do their due diligence. If the algorithm makes a mistake, it is harder to explain"* (P14). In the LEG context, P05 describes challenges with AI-generated risk scores for recidivism replacing expert testimony:

"It used to be that you would work with probation officers and other experts who would say, 'Hey! Given this person's age, given their criminal history background, given the age at which they started being in the system, it's more or less likely they will re-offend.' And now, you get a score. For a defense attorney, it's like, 'Well, what am I supposed to do with that score?' If this were an expert testifying, I would have the right to cross-examine them. We would have the ability to question their expertise... We are essentially allowing these AI systems to provide expert testimony without any of the protections that we normally afford to expert testimony." (P05)

P09 clarified that decision-makers are *"not new to that type of tension where you're confused about something and asking for help"* as MED decision-makers often consult their peers and are used to seeking second opinions (e.g., clinicians frequently come together to discuss complex cases during grand rounds). Current AI tools for decision-making often only provide one-shot recommendations, which contrasts with established practices of dialogue and clarification that allow decision-makers multiple opportunities to ask questions. Comparatively, participants spoke more positively about generative chatbots, as these tools enable back-and-forth discussion (P03, P09, P11).

Faulty beliefs in model capabilities. Adoption choices are also shaped by what decision-makers believe AI tools can do—which may not always be accurate. Unclear model limitations can lead to misguided assumptions about model capabilities and adoption benefits. As one participant noted, AI is sometimes given *"an unnecessary and undeserved veneer of neutrality and objectivity,"* even though models and their training data are *"not inherently objective or neutral, and reflect our own biases"* (P05). The challenge of accurately assessing model quality is further compounded by the rapid proliferation of models across both research and commercial domains. As noted by P08, models are being *"created by all of these different people,"* making evaluation difficult. Moreover, limited clinical or professional engagement with modern machine learning methods (e.g., deep learning, neural networks) leaves some decision-makers struggling to *"judge what's a really good model"* (P08). As a result, they

may base adoption choices on faulty beliefs and overconfidence in AI capabilities.

4.3 Consequences for the Decision-maker

When choosing whether to adopt AI tools, decision-makers also weigh potential consequences—including additional **legal, professional, and reputational liabilities**—against their current workflows. In particular, decision-makers may consider the risk of potential lawsuits, loss of professional licences, or damaging credibility, citing concerns such as “*the reputational risk of being wrong*” (P12). Legal liability remains a complex consideration as “*laws have still not evolved*” to address decisions involving AI models (P09). In MED, P10 emphasizes that “*regulating AI models is hard and it all moves at a pace that the FDA does not*”. Similarly, in LEG contexts, although judges and court systems have begun to issue guidance rules on the use of AI (e.g., the guidance for the use of generative AI from the state bar of California mentioned by P06), regulation is still largely “*catching up*” to actual AI use in practice (P05).

In addition, **uncertainties about how these consequences apply** may also discourage adoption. For instance, clinicians in MED ask:

“Who is liable? Is it the AI? Is it the company or the programmer that made it? That then a health system purchased? Was [the model] written in-house? Is it the doctor who, at the end of the whole process, is ultimately responsible?” (P14)

Beyond legal liability, there are also concerns about political and professional risks, among other liabilities. P05 described how judges may be hesitant to rely on AI tools for fear that incorrect decisions could become a “*political liability*” if they later run for office (P05). While some AI vendors have attempted to mitigate these concerns by offering indemnification clauses—e.g., compensating law firms for losses caused by chatbot errors—the individual professional often remains ultimately accountable: “*to the client, the lawyer is [still] the sole responsible party for any mistakes*” (P06). As P05 summarized, adopting AI tools is often seen as being “*at your own peril*”.

Beyond legal and reputational consequences, decision-makers also consider practical implications for their daily work, such as **workload changes**. Many participants mentioned poor integration with existing workflows as a significant barrier to adoption (P10, P13, P14). Especially in high-stress fields where burnout is prevalent, any additional effort required to incorporate AI tools can be prohibitive (P09). In contrast, participants expressed more optimism toward AI adoption in “*low-stakes*” scenarios (i.e., contexts where there are fewer potential negative consequences and the impact on workload is more clearly beneficial). For example, P11 noted that “*administrative simplification and cutting out processes that are low-stakes and do not have risk for bias are a lot easier to contemplate using AI for than having an agent-based model make decisions that impact patients.*” Similarly, P08 described how most AI models currently used in practice address “*low-hanging fruit and low-stakes problems such as resource utilization or administrative work.*”

In these cases, decision-makers can leverage the processing capability of AI tools “*to streamline their processes to make their work a little bit more efficient*” (P15).

4.4 Perceived Implications for Other Stakeholders

When considering AI adoption, decision-makers also weigh the broader impact of AI and the implications for other stakeholders involved in how AI systems are deployed and used. These stakeholders’ goals may diverge from those of the decision-maker. One key stakeholder is the **organization** to which the decision-maker belongs. For instance, in JRN, news agencies “*obviously have different agendas*” from journalists and typically prioritize profits P02, whereas many journalists value producing high-quality content over maximizing output. Nonetheless, organizational objectives such as cost reduction and efficiency gains can significantly influence decision-makers’ adoption choices.

Decision-subjects (e.g., patients in MED settings or clients in LEG contexts) are another crucial stakeholder group considered by decision-makers. While AI tools can offer benefits to decision-makers and their organizations, participants noted potential harms to decision-subjects. For example, AI systems may fail to identify contextually relevant factors that a human might catch (P05), or introduce biases and stereotyping in decision-making (P04, P06, P09, P10, P12, P15). These risks may make decision-makers hesitant to adopt AI tools, especially when decision-makers perceive risks to vulnerable populations.

5 Comparative Analysis of AI Adoption

We use the factors outlined in Section 4 to develop an AI adoption sheet (Figure 1) that helps analyze decision-makers’ adoption choices around adopting AI for specific use cases. We define these use cases as unique contexts where AI supports decision-making or task completion. Importantly, decision-maker adoption is often context-dependent, varying not only across different use cases, but also for similar use cases in different domains.

To identify use cases of interest, we first coded 34 distinct use cases mentioned by our participants (detailed in the Appendix). Most use cases are domain-specific, targeting specific tasks or objectives (e.g., recidivism prediction in LEG or MED diagnostics through radiology reads). However, two use cases appeared in three domains each: **e-discovery** (mentioned by nine participants in JRN, LEG, PHW) and **tailoring communications** (mentioned by four participants in JRN, MED, PHW). We now examine both case studies through a cross-domain lens to illustrate how our AI adoption sheet helps systematically elucidate differences in decision-maker adoption.

5.1 E-Discovery for Research and Analysis

Decision-makers increasingly use AI tools for *e-discovery*, which refers to gathering, organizing, and analyzing relevant digital data for downstream goals. For example, journalists investigating a story often need to draw on various sources of data (e.g., interviews, public records) (Stray 2021; Broussard

USE CASE E-discovery	Journalism	Legal	Medicine	Public Sector
Decision-maker Background	+	-	-	-
Model Perception	+	+	-	+
Consequences for the Decision-maker	~	-	-	~
Perceived Implications for Other Stakeholders	+	+	-	-

(a) E-Discovery for Research and Analysis

USE CASE Tailoring Communications	Journalism	Legal	Medicine	Public Sector
Decision-maker Background	+	-	-	+
Model Perception	~	-	+	+
Consequences for the Decision-maker	~	~	~	~
Perceived Implications for Other Stakeholders	-	~	+	-

(b) Tailoring Communications to Audiences of Different Backgrounds

Figure 4: An overview of two cross-domain case studies analyzed through our AI adoption sheet. For each domain, a green plus sign (+) indicates that a factor is likely to influence decision-makers to adopt AI tools, a red minus sign (-) indicates that a factor is likely to sway decision-makers NOT to adopt AI tools, and a tilde (~) indicates that the factor likely does not strongly affect a decision-maker’s adoption choice.

et al. 2019; Quinonez and Meij 2024). All of our JRN participants described AI tools for e-discovery to be extremely helpful in streamlining the reporting process (P01, P02, P03, P04). In LEG contexts, discovery is a long-standing procedure that comprises the initial phase of litigation, requiring attorneys and paralegals to sift through case materials to identify pertinent documents and evidence. Recently, AI tools have allowed attorneys to more efficiently identify relevant documents in the increasingly digitalized landscape through e-discovery (P05, P06) (Pai et al. 2023; Nagineni 2024; Yang et al. 2024a). PHW decision-makers similarly draw on community-level data to help shape effective policies, and participants described the use of AI tools for the analysis of these data (P13, P15, P16) (Shang et al. 2024; Yang et al. 2024b). In MED settings, a key part of clinicians’ workflow involves reviewing patient histories to tailor treatment plans and provide the highest quality of care for each patient. AI for e-discovery could substantially reduce their workload; however, our MED participants noted that these

types of tools are seldom adopted in practice (P08, P12).¹ We summarize our analysis in Figure 4a, where we illustrate why e-discovery AI tools are more widely adopted by decision-makers in some domains (e.g., in JRN, LEG, PHW) compared to others (i.e., in MED).

Factor 1: Decision-maker background. In fields requiring high specialization, such as LEG, MED, and PHW, decision-makers often have a greater sense of confidence in their own judgment. Such confidence may make them less inclined to rely on support from AI tools, especially if the tools might lack relevant context (P05, P09). For example, an attorney may believe “AI tools have a hard time understanding the nuance and complex context involved” in case documents and thus prefer to go through discovery on their own (P06). P05 notes that although AI tools for e-discovery can be extremely beneficial, “there’s a trade-off. Being the associate who knew the entire case and read every document

¹A large-scale survey by PWC studying AI and robotics in healthcare has also not identified the use of AI tools for e-discovery in MED (PWC 2017).

was really valuable and obviously, there are times when a machine learning algorithm is not able to correctly predict what would have been relevant whereas a lawyer would have noticed something and pulled that document.” We also observe this in MED, where clinicians may prioritize their professional experience over insights provided by AI (P08, P12). Even in PHW, case workers often have extensive professional experience with the communities they serve, allowing them to process more nuanced context than AI tools can provide (P14). In contrast, decision-makers in JRN may tackle a wide range of topics, including those outside their immediate expertise, which may make them more likely to embrace AI tools for e-discovery (P02, P03).

Factor 2: Model perception. E-discovery tools in JRN and LEG settings are typically used for researching information as decision-makers believe they can simplify and streamline research processes effectively (P03, P05, P06, (Warren et al. 2024)). Similarly, decision-makers believe e-discovery tools in PHW can be used to support existing data analysis frameworks (e.g., generating spreadsheet formulas). As P15 explained: “If I ask an AI tool for a formula to do something in sheets and it gives me an answer that’s not true, it’s easy for me to know that it’s not true because when I just go in, I type in the formula and it doesn’t work”. In these settings, decision-makers may perceive models to be more trustworthy as it is easier to identify when AI outputs are incorrect. In contrast, participants in MED domains tend to express significant skepticism towards the use of AI tools and algorithms in clinical practice (P10) as it remains “unclear what their added value is and whether there is really consensus that these models are safe to use”(P08). P09 admits it feels “a bit off when you have the human depending on the accuracy of the model and not on the expertise of humans.” Especially as there are a multitude of established and validated resources in MED (e.g., classical risk scores), clinicians may be unlikely to adopt AI tools for e-discovery.

Factor 3: Consequences for the decision-maker. The stakes of using e-discovery tools vary across domains. For example, attorneys may face serious legal repercussions if e-discovery tools cause pertinent information to be missed or lead the attorney to form biased opinions, negatively impacting trial outcomes. Improper handling of sensitive data may also have legal ramifications for the attorney considering “chats with AI systems are not confidential” (P06) and would “be a breach of duties as a lawyer to protect client confidential information,” violating attorney-client privilege without informed consent (P05). In MED, similar concerns have arisen regarding incorporating e-discovery tools, though the full extent of potential consequences remains unclear. As noted by P10, malpractice liability is a key consideration among clinicians:

“All the clinicians have malpractice insurance. Do the algorithms? Do they need to? I think all of these are hard questions. If I was liable for what happened to my patient and I have very expensive insurance because of that, I don’t want to be told what to do unless they’re going to take the liability away.” (P10)

In high-stakes domains such as LEG and MED—where risks

are more pronounced—decision-makers may be more hesitant to adopt tools that could expose them to additional liability. In comparison, mistakes as a result of using AI tools during the research step in both JRN and PHW may only result in a mild reprimand.

Factor 4: Perceived implications for other stakeholders. Since e-discovery tools can make decision-makers more efficient, other stakeholders may prefer decision-makers to adopt these tools. For example, P05 and P06 noted that e-discovery tools powered by generative AI, when used appropriately, can help not only lawyers, but also “clients because it can make what lawyers do cheaper” as clients are billed based on the hours spent on their case (P05). Some law firms may also encourage the usage of these tools so that lawyers can take on more clients. Similarly, news companies that journalists work for may also urge the use of e-discovery tools to help maximize content output, and thus business profits as well (P02, P03). However, in MED, while some hospitals may promote the adoption of AI for e-discovery to increase the number of patients that a clinician sees each day, clinicians tend to feel that they must take the patient into account above everything else. To this end, P08 emphasized that AI outputs, which oftentimes “objectively seem to be the best, may subjectively be worse for the patient” and may cause more harm than good if adopted. Analogously in PHW, case workers often work for non-profits or government agencies that have “limited capacity to maintain and use more complex models” (P14), which may result in a limited number of tools available to decision-makers.

5.2 Tailoring Communication to Audiences of Differing Backgrounds

Effective communication is crucial for successful interactions across numerous domains and is often necessary to communicate between individuals from different backgrounds (e.g., educational, cultural, linguistic). Tailoring communication to audiences of differing backgrounds is a complex task that has seen an increase in the use of AI tools (Abdelali, Bennoudi et al. 2023; Algouzi and Alzubi 2023; Ayvazyan, Hao, and Pym 2024; Dwivedi et al. 2024; Lee 2024). Our participants also noted the use of AI tools for relevant applications, including language translation (P01, P03), simplifying jargon (P09, P12), and even adjusting the tone and style of writing to better align with specific cultural contexts (P15, P16). Despite their potential, these AI tools are not always adopted by decision-makers. While this use case emerged in interviews across JRN, MED, and PHW, it was absent in LEG.² Using our factors, we investigate how AI tools are used to tailor communications across domains (summarized in Figure 4b). We focus our discussion on Factors 1, 2, and 4 as they have the greatest inter-domain differences. For conciseness, we omit discussion of Factor 3 as inter-domain differences are less distinct.

Factor 1: Decision-maker background. Emphasis in some domains prioritizing human-to-human interactions may discourage the adoption of AI for tailoring commu-

²It similarly did not appear across surveys in LEG conducted by Warren et al. (2024).

nications (P05, P06, P08, P09). Especially in LEG where “lawyers are the world’s most highly paid rhetoricians” (P05), direct communications with the decision-maker may be highly valued. Similarly, AI tools can help streamline communications in MED, but “shared decision-making between a doctor and their patient—which has to be customized to what the patient feels is important at that time”—can never be replaced by AI due to the subjective nature of MED (P08). Alternatively, fields such as PHW may be “chronically underfunded, underresourced, understaffed, and overworked”, so a desire to “be more efficient and do more with less” may encourage decision-makers to adopt AI tools to avoid burnout (P15).

Factor 2: Model perception. Depending on the specific needs of a decision-maker, they may perceive model performance and reliability differently. Although some decision-makers may be inclined to use AI tools for translation as many languages are available, others may be reluctant to adopt such tools because they feel “some languages are not yet working well” (P16) or believe that AI generates low-quality outputs (P04, (Thompson et al. 2024)). Participants echo the sentiment found in Clio (2023), which highlights that LEG professionals are interested yet cautious of AI due to the risk of hallucinations (P05, P06). Similarly, in MED, both our participants (P09, P11) and the surveys by Association (2025) show that physician sentiment toward enthusiasm for AI translation services trends positive, with a 3% increase from 2023 to 2024 (Association 2025). P15 explains this with a concrete example: When prompting an AI model for a vaccine analogy, model outputs “tailoring for the Latino and Hispanic communities gave an example of a soccer game” but when prompting for anecdotes “tailored to the Black and African American communities, the model gave me a bulletproof vest analogy which, of course, is hugely problematic.” However, in a recent repeat of this experiment, P15 found the model had improved and no longer performed in this problematic way.

Factor 4: Perceived implications for other stakeholders. AI tools can enable decision-makers to tailor communications to a variety of backgrounds. In MED, while “the best clinicians are the ones who personalize their decisions around the patient,” patients must be able to fully understand any health considerations to properly “come up with a shared decision” with their clinician (P08). As such, our participants spoke positively about AI tools that can translate medical text into “a reading format that fits the patient and their education level” (P09) to “help patients or caregivers when they have to make personal, value-based decisions by turning something of high complexity into something understandable for them” (P12). Patients have even been found to prefer AI responses to clinician-to-patient messages (Kim et al. 2024; Tai-Seale et al. 2024). However, when decision-makers believe that AI tools fail to account for varying cultural backgrounds and the impact of harmful speech, they may be less inclined to adopt AI tools. This is because the outputs of AI models “become representative of a country and representative of a culture”, and models trained on low-quality data can perpetuate harmful stereotypes (P04). Additionally, P15 notes that AI use can have a “broader impact. If

we continue to use language that is harmful to communities, if we use language that is inaccurate, if we use tools that provide information that is misleading, or that turns people off to the information because it is not presented in a compassionate way. I think all of that will have a negative impact down the line” such as “contribute to inequities and contribute to widening disparities” (P15).

6 Discussion

6.1 Recommendation for Better AI Adoption

While decision-makers ultimately choose whether to adopt AI tools, their choices are shaped by many other actors—including model developers, policymakers, and researchers. Building on our findings, we offer four recommendations to support responsible AI adoption in practice.

(1) Integrate Decision-Maker Perspectives in AI Development. We recommend *starting with the decision-maker and considering whether they are likely to adopt an AI tool* before developing and deploying AI tools for specific use cases. We present our AI adoption sheet in Figure 1, summarizing the four high-level factors that we identified. We also provide examples of how this AI adoption sheet can be applied to the aforementioned case studies in the Appendix. Prior to deploying new AI tools, we encourage model developers, policymakers, and researchers to utilize this AI adoption sheet to proactively assess how AI tools may be adopted by decision-makers in practice.

(2) Bolster Decision-maker Education on AI. Increasing AI literacy can mitigate the influence of decision-maker background and model misperception on adoption choices (P07, P08, P16). We encourage decision-makers to engage with emerging resources on AI literacy (Ng et al. 2021; Chiu et al. 2024). We also recommend that organizations developing AI tools concurrently invest in literacy initiatives in collaboration with AI researchers and model designers. In particular, we encourage future literacy work to focus not only on decision-maker adoption choices, but also on the relationship between adoption and reliance to promote responsible use of AI tools by decision-makers.

(3) Develop Infrastructure to Integrate AI into Existing Workflows. Oftent, the additional workload required to use AI tools not integrated into decision-maker workflows may be prohibitive to adoption. Moreover, AI systems often demand substantial resources—such as data storage—that many organizations, particularly those lacking a clear AI adoption strategy, struggle to accommodate (P01, P06, P10, P13, P14). These challenges typically place an additional burden on the decision-maker, further deterring adoption. We recommend that organizations adopt a more deliberate approach to AI procurement, ensuring they have the necessary infrastructure in place. To support decision-maker adoption, deployed tools should align with existing workflows rather than requiring substantial procedural changes (Silva et al. 2025).

(4) Provide Resources to Decision-makers for Navigating Applicable Policies. The consequences decision-makers must weigh when adopting AI tools remain poorly defined (P10, P11). Our interviews suggest that the current regu-

latory landscape is inadequate, with one participant noting a “lack [of] the right regulatory framework” to ensure responsible AI use (P05). The reactive nature of AI policy has left decision-makers struggling to “navigate the patchwork regulatory approach that currently exists” (P05). Policymakers should interface with existing research on AI governance (Schiff et al. 2020; Stix 2021; Schiff 2022) and collaborate with researchers, model designers, and decision-makers to craft more coherent and proactive AI policies. However, until such policies are implemented, it is essential that organizations offer practical resources to their members to support navigation of the existing regulatory environment.

6.2 Comparison of Findings to Prior Work

Comparison to organizational factors. As highlighted in Section 2.2, researchers have extensively studied AI adoption at the organizational level, identifying factors such as issues with data quality and availability, lack of data governance or regulation, and goals related to cost or time efficiency. These organizational-level factors parallel some of our decision-maker-centric factors—for instance, organizational concerns about data quality or availability mirror decision-maker concerns about model performance (*Model perception*). However, prior work studying organizational adoption tends to overlook the nuanced and multifaceted consequences considered by decision-makers, which include time trade-offs, various types of liabilities, and implications for other stakeholders that may not align with organizational priorities. Although organizational factors touch on some of these issues, they typically focus narrowly on organizational goals. Examining adoption at both the decision-maker and organizational levels highlights overlapping factors that are important to consider when evaluating AI tools. However, considering AI adoption from the perspective of decision-makers may offer a more comprehensive understanding of the factors shaping real-world adoption.

Comparison to previously identified decision-maker factors. While prior work has rarely studied adoption at the level of the individual decision-maker, our findings align with and extend insights from two prior works (Fang et al. 2023; Hameed et al. 2023). Fang et al. (2023) emphasizes that distinctions based on professional roles—analogue to *Decision-maker background*—are important to understanding decision-maker adoption in LEG. However, their findings focus solely on how occupational roles shape adoption attitudes and do not cover any other factors. Hameed et al. (2023) validates seven predefined constructs to influence adoption behaviors (performance expectancy, effort expectancy, personal innovativeness in information technology, task complexity, technological characteristics, initial trust, and social influence) that also align with our identified factors. However, while these studies find factors to always be significant, our case studies in Section 5 show that these factors vary in importance across domains and use cases. By offering broader categories with more granular prompts, our AI adoption sheet helps organize and simplify these constructs, encouraging more systematic future research into what shapes AI adoption choices. Together, these contribu-

tions provide a structured, practical tool for future research and practice.

6.3 Limitations and Future Work

Our study has several limitations, which motivate directions for future work. We acknowledge that our participant sample size and the set of domains are limited and do not capture the full spectrum of real-world applications. Additionally, our participants were all based in the United States, which may limit the generalizability of our findings to other domains and contexts. As such, the four key factors we identified in Section 4 may not be exhaustive. We encourage future work to build on our methods and evaluate a larger sample of decision-makers in more domains. Additionally, as our recruitment protocol required that study participants have prior experience with AI tools in their domain, our interviewees may be biased towards the adoption of AI tools in practice. Participants may have also had varying degrees of comfort when discussing some use cases and factors over others, which may have shaped the responses we received. As our interviews were open-ended, discussions primarily focused on an *individual* decision-maker, which is consistent with most prior work. However, we encourage future work to also consider factors influencing adoption by groups of decision-makers. Further, we acknowledge that decision-makers’ perceptions and the capabilities of AI models are rapidly evolving. Our interviews, conducted between October 2023 and October 2024, reflect perceptions and model capabilities relevant during this time frame. Finally, we focused on *inter-domain* differences; further study is needed to explore *intra-domain* differences between decision-makers and how they influence AI adoption.

7 Conclusion

AI adoption depends on the decision-maker actively incorporating AI tools into their workflows. To identify common factors across domains that influence decision-makers’ use of AI, we conducted 16 semi-structured interviews with decision-makers across four domains. Through our cross-domain analysis, we identify four key factors that impact decision-maker adoption. We apply this sheet to two case studies of real-world use cases, illustrating why AI-powered tools are more widely adopted by decision-makers in some domains and not others. We believe that the AI adoption sheet has many uses; it can support researchers considering the adoption of AI in new domains, model developers creating new AI tools with the end user in mind, and organizations identifying appropriate AI tools suited for decision-maker adoption. In closing, we hope that this work not only encourages further analyses focused on the decision-maker but also inspires future research to look beyond a single use case or domain.

Acknowledgements

Heidari and Yu acknowledges support from NSF (IIS2040929 and IIS2229881). Any opinions, findings, conclusions, or recommendations expressed in this material

are those of the authors and do not reflect the views of NSF or other funding agencies.

References

- Abdelali, Z.; Bennoudi, H.; et al. 2023. AI vs. Human Translators: Navigating the Complex World of Religious Texts and Cultural Sensitivity. *International Journal of Linguistics, Literature and Translation*, 6(11): 173–182.
- Algouzi, S.; and Alzubi, A. A. F. 2023. The study of AI-mediated communication and socio-cultural language-related variables: Gmail reply suggestions. *Applied Artificial Intelligence*, 37(1): 2175114.
- Amiri, P.; and Karahanna, E. 2022. Chatbot use cases in the Covid-19 public health response. *Journal of the American Medical Informatics Association*, 29(5): 1000–1010.
- Association, A. M. 2025. AMA Augmented Intelligence Research: Physician sentiments around the use of AI in healthcare. Technical report, American Medical Association.
- Ayvazyan, N.; Hao, Y.; and Pym, A. 2024. Things to Do in the Translation Class When Technologies Change: The Case of Generative AI. In *New Advances in Translation Technology: Applications and Pedagogy*, 219–238. Springer.
- Bansal, G.; Nushi, B.; Kamar, E.; Lasecki, W. S.; Weld, D. S.; and Horvitz, E. 2019. Beyond accuracy: The role of mental models in human-AI team performance. In *Proceedings of the AAAI conference on human computation and crowdsourcing*, volume 7, 2–11.
- Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. 2021. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, 1–16.
- Başar, E. E.; and Erkul, A. K. 2024. Factors Affecting the Attitude of Medical Doctors in Türkiye towards Using Artificial Intelligence Applications in Healthcare Services. *Bezmialem Science*.
- Bedué, P.; and Fritzsche, A. 2022. Can we trust AI? An empirical investigation of trust requirements and guide to successful AI adoption. *Journal of Enterprise Information Management*, 35(2): 530–549.
- Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2): 77–101.
- Braun, V.; and Clarke, V. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health*, 11(4): 589–597.
- Braun, V.; and Clarke, V. 2021. Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. *Counselling and psychotherapy research*, 21(1): 37–47.
- Broussard, M.; Diakopoulos, N.; Guzman, A. L.; Abebe, R.; Dupagne, M.; and Chuan, C.-H. 2019. Artificial intelligence and journalism. *Journalism & mass communication quarterly*, 96(3): 673–695.
- Cai, C. J.; Winter, S.; Steiner, D.; Wilcox, L.; and Terry, M. 2019. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).
- Campion, A.; Gasco-Hernandez, M.; Jankin Mikhaylov, S.; and Esteve, M. 2022. Overcoming the challenges of collaboratively adopting artificial intelligence in the public sector. *Social Science Computer Review*, 40(2): 462–477.
- Chatterjee, S.; Rana, N. P.; Dwivedi, Y. K.; and Baabdullah, A. M. 2021. Understanding AI adoption in manufacturing and production firms using an integrated TAM-TOE model. *Technological Forecasting and Social Change*, 170: 120880.
- Chen, V.; Liao, Q. V.; Wortman Vaughan, J.; and Bansal, G. 2023. Understanding the role of human intuition on reliance in human-AI decision-making with explanations. *Proceedings of the ACM on Human-computer Interaction*, 7(CSCW2): 1–32.
- Chiu, T. K.; Ahmad, Z.; Ismailov, M.; and Sanusi, I. T. 2024. What are artificial intelligence literacy and competency? A comprehensive framework to support them. *Computers and Education Open*, 6: 100171.
- Clio. 2023. Legal Trends Report. Technical report, Clio.
- Contini, F. 2024. Unboxing Generative AI for the Legal Professions: Functions, Impacts and Governance. In *IJCA*, volume 15, 1. HeinOnline.
- Cubric, M. 2020. Drivers, barriers and social considerations for AI adoption in business and management: A tertiary study. *Technology in Society*, 62: 101257.
- Davenport, T. H.; and Glaser, J. P. 2022. Factors governing the adoption of artificial intelligence in healthcare providers. *Discover Health Systems*, 1(1): 4.
- Dietvorst, B. J.; Simmons, J. P.; and Massey, C. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of experimental psychology: General*, 144(1): 114.
- Dwivedi, A.; Sahu, S.; Srivastava, A.; et al. 2024. AI-Powered Real-Time Text Editor with Multilingual Translation and Speech Recognition. In *2024 IEEE 16th International Conference on Computational Intelligence and Communication Networks (CICN)*, 1335–1340. IEEE.
- Dwivedi, Y. K.; Hughes, L.; Ismagilova, E.; Aarts, G.; Coombs, C.; Crick, T.; Duan, Y.; Dwivedi, R.; Edwards, J.; Eirug, A.; et al. 2021. Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International journal of information management*, 57: 101994.
- Fang, C.; Wilkenfeld, J. N.; Navick, N.; and Gibbs, J. L. 2023. "AI Am Here to Represent You": Understanding How Institutional Logics Shape Attitudes Toward Intelligent Technologies in Legal Work. *Management Communication Quarterly*, 37(4): 941–970.
- Goodman, L. A. 1961. Snowball sampling. *The annals of mathematical statistics*, 148–170.

- Hameed, B. Z.; Naik, N.; Ibrahim, S.; Tatkar, N. S.; Shah, M. J.; Prasad, D.; Hegde, P.; Chlosta, P.; Rai, B. P.; and Somani, B. K. 2023. Breaking barriers: Unveiling factors influencing the adoption of artificial intelligence by healthcare providers. *Big Data and Cognitive Computing*, 7(2): 105.
- Harvey, H. B.; and Gowda, V. 2021. Regulatory issues and challenges to artificial intelligence adoption. *Radiologic Clinics*, 59(6): 1075–1083.
- Henry, K. E.; Kornfield, R.; Sridharan, A.; Linton, R. C.; Groh, C.; Wang, T.; Wu, A.; Mutlu, B.; and Saria, S. 2022. Human–machine teaming is key to AI adoption: clinicians’ experiences with a deployed machine learning system. *NPJ digital medicine*, 5(1): 97.
- Jo, E.; Epstein, D. A.; Jung, H.; and Kim, Y.-H. 2023. Understanding the Benefits and Challenges of Deploying Conversational AI Leveraging Large Language Models for Public Health Intervention. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.
- Kawakami, A.; Coston, A.; Heidari, H.; Holstein, K.; and Zhu, H. 2024. Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions Around AI Design and Use. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW2): 1–24.
- Kawakami, A.; Sivaraman, V.; Cheng, H.-F.; Stapleton, L.; Cheng, Y.; Qing, D.; Perer, A.; Wu, Z. S.; Zhu, H.; and Holstein, K. 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*, 1–18.
- Kim, J.; Chen, M. L.; Rezaei, S. J.; Liang, A. S.; Seav, S. M.; Onyeka, S.; Lee, J. J.; Vedak, S. C.; Mui, D.; Lal, R. A.; et al. 2024. Perspectives on Artificial Intelligence–Generated Responses to Patient Messages. *JAMA Network Open*, 7(10): e2438535–e2438535.
- Lee, M. K.; Grgić-Hlača, N.; Tschantz, M. C.; Binns, R.; Weller, A.; Carney, M.; and Inkpen, K. 2020. Human-Centered Approaches to Fair and Responsible AI. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI EA ’20, 1–8. New York, NY, USA: Association for Computing Machinery. ISBN 9781450368193.
- Lee, T. K. 2024. Artificial intelligence and posthumanist translation: ChatGPT versus the translator. *Applied Linguistics Review*, 15(6): 2351–2372.
- Logg, J. M.; Minson, J. A.; and Moore, D. A. 2019. Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151: 90–103.
- Madan, R.; and Ashok, M. 2023. AI adoption and diffusion in public administration: A systematic literature review and future research agenda. *Government Information Quarterly*, 40(1): 101774.
- Merhi, M. I. 2023. An evaluation of the critical success factors impacting artificial intelligence implementation. *International Journal of Information Management*, 69: 102545.
- Nagineni, P. 2024. The Rise of AI in eDiscovery: How Machine Learning is Revolutionizing Legal Data Processing. *International Journal of Computer Engineering and Tehnology (IJCET)*, 15(5): 329–341.
- Neumann, O.; Guirguis, K.; and Steiner, R. 2024. Exploring artificial intelligence adoption in public organizations: a comparative case study. *Public Management Review*, 26(1): 114–141.
- Ng, D. T. K.; Leung, J. K. L.; Chu, S. K. W.; and Qiao, M. S. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2: 100041.
- Pai, S.; Lahiri, S.; Kumar, U.; Baksi, K.; Soba, E.; Suesserman, M.; Pudota, N.; Foster, J.; Bowen, E.; and Bhat-tacharya, S. 2023. Exploration of Open Large Language Models for eDiscovery. In *Proceedings of the Natural Legal Language Processing Workshop 2023*, 166–177.
- Park, S. Y.; Kuo, P.-Y.; Barbarin, A.; Kaziunas, E.; Chow, A.; Singh, K.; Wilcox, L.; and Lasecki, W. S. 2019. Identifying challenges and opportunities in human-AI collaboration in healthcare. In *Companion Publication of the 2019 Conference on Computer Supported Cooperative Work and Social Computing*, 506–510.
- Parker, C.; Scott, S.; and Geddes, A. 2019. Snowball sampling. *SAGE research methods foundations*.
- Passi, S.; and Vorvoreanu, M. 2022. Overreliance on AI literature review. *Microsoft Research*, 339: 340.
- Polisetty, A.; Chakraborty, D.; Kar, A. K.; and Pahari, S. 2024. What determines AI adoption in companies? Mixed-method evidence. *Journal of Computer Information Systems*, 64(3): 370–387.
- Poon, E. G.; Lemak, C. H.; Rojas, J. C.; Guptill, J.; and Classen, D. 2025. Adoption of artificial intelligence in healthcare: survey of health system priorities, successes, and challenges. *Journal of the American Medical Informatics Association*, ocaf065.
- PWC. 2017. Why AI and Robotics will define New Health. Technical report, PWC.
- Quinonez, C.; and Meij, E. 2024. A new era of AI-assisted journalism at Bloomberg. *AI Magazine*.
- Regona, M.; Yigitcanlar, T.; Xia, B.; and Li, R. Y. M. 2022. Opportunities and adoption challenges of AI in the construction industry: A PRISMA review. *Journal of open innovation: technology, market, and complexity*, 8(1): 45.
- Reverberi, C.; Rigon, T.; Solari, A.; Hassan, C.; Cherubini, P.; and Cherubini, A. 2022. Experimental evidence of effective human–AI collaboration in medical decision-making. *Scientific reports*, 12(1): 14952.
- Schelble, B. G.; Lopez, J.; Textor, C.; Zhang, R.; McNeese, N. J.; Pak, R.; and Freeman, G. 2024. Towards ethical AI: Empirically investigating dimensions of AI ethics, trust repair, and performance in human-AI teaming. *Human Factors*, 66(4): 1037–1055.

- Schemmer, M.; Kuehl, N.; Benz, C.; Bartos, A.; and Satzger, G. 2023. Appropriate reliance on AI advice: Conceptualization and the effect of explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*, 410–422.
- Schiff, D. 2022. Education for AI, not AI for education: The role of education and ethics in national AI policy strategies. *International Journal of Artificial Intelligence in Education*, 32(3): 527–563.
- Schiff, D.; Biddle, J.; Borenstein, J.; and Laas, K. 2020. What's next for ai ethics, policy, and governance? a global overview. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 153–158.
- Schoeffer, J.; De-Arteaga, M.; and Kuehl, N. 2024. Explanations, fairness, and appropriate reliance in human-AI decision-making. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Sethi, T.; Kalia, A.; Sharma, A.; and Nagori, A. 2020. Interpretable artificial intelligence: Closing the adoption gap in healthcare. In *Artificial Intelligence in Precision Health*, 3–29. Elsevier.
- Shang, T.; Yang, S.; He, W.; Zhai, T.; Li, D.; Hou, B.; Chen, T.; Moore, J. H.; Ritchie, M. D.; and Shen, L. 2024. Leveraging Social Determinants of Health in Alzheimer's Research Using LLM-Augmented Literature Mining and Knowledge Graphs. *arXiv preprint arXiv:2410.09080*.
- Silva, E.; Johnson, N.; Dotan, R.; Eslami, M.; Heidari, H.; and Schwanke, B. 2025. Procuring Public-Sector AI: Guidance for Local Governments. Available at SSRN 5227802.
- Silva, Í.; Marinho, L.; Said, A.; and Willemsen, M. C. 2024. Leveraging ChatGPT for Automated Human-centered Explanations in Recommender Systems. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, 597–608.
- Stix, C. 2021. Actionable principles for artificial intelligence policy: three pathways. *Science and engineering ethics*, 27(1): 15.
- Stray, J. 2021. Making artificial intelligence work for investigative journalism. *Algorithms, Automation, and News*, 97–118.
- Sun, T. Q.; and Medaglia, R. 2019. Mapping the challenges of Artificial Intelligence in the public sector: Evidence from public healthcare. *Government information quarterly*, 36(2): 368–383.
- Tai-Seale, M.; Baxter, S. L.; Vaida, F.; Walker, A.; Sitapati, A. M.; Osborne, C.; Diaz, J.; Desai, N.; Webb, S.; Polston, G.; et al. 2024. AI-generated draft replies integrated into health records and physicians' electronic communication. *JAMA Network Open*, 7(4): e246565–e246565.
- Thieme, A.; Hanratty, M.; Lyons, M.; Palacios, J.; Marques, R. F.; Morrison, C.; and Doherty, G. 2023. Designing Human-centered AI for Mental Health: Developing Clinically Relevant Applications for Online CBT Treatment. *ACM Trans. Comput.-Hum. Interact.*, 30(2).
- Thompson, B.; Dhaliwal, M.; Frisch, P.; Domhan, T.; and Federico, M. 2024. A shocking amount of the web is machine translated: Insights from multi-way parallelism. In *Findings of the Association for Computational Linguistics ACL 2024*, 1763–1775.
- Uren, V.; and Edwards, J. S. 2023. Technology readiness and the organizational journey towards AI adoption: An empirical study. *International Journal of Information Management*, 68: 102588.
- Verma, D. 2024. Impact of Artificial Intelligence on Journalism: A Comprehensive Review of AI in Journalism. *Journal of Communication and Management*, 3(02): 150–156.
- Warren, Z.; Abbott, M.; Leach, L.; Post, C.; Seemer, S.; Strunsky, A.; and Wirth, G. 2024. 2024 Generative AI in Professional Services: Perceptions, Usage and Impact on the Future of Work. Technical report, Thomson Reuters Institute.
- Wu, K.; Wu, E.; Theodorou, B.; Liang, W.; Mack, C.; Glass, L.; Sun, J.; and Zou, J. 2024. Characterizing the clinical adoption of medical AI devices through US insurance claims. *NEJM AI*, 1(1): AIoa2300030.
- Yang, E.; Omrani, R.; Curtin, E.; Emory, T.; Gray, L.; Pickens, J.; Reff, N.; Traylor, C.; Underwood, S.; Lewis, D. D.; et al. 2024a. Beyond the Bar: Generative AI as a Transformative Component in Legal Document Review. In *2024 IEEE International Conference on Big Data (Big-Data)*, 4779–4788. IEEE.
- Yang, J.; Walker, K. C.; Bekar-Cesaretti, A. A.; Hao, B.; Bhadelia, N.; Joseph-McCarthy, D.; and Paschalidis, I. C. 2024b. Automating biomedical literature review for rapid drug discovery: Leveraging GPT-4 to expedite pandemic response. *International Journal of Medical Informatics*, 105500.
- Yang, Q.; Steinfeld, A.; and Zimmerman, J. 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*, CHI '19, 1–11. New York, NY, USA: Association for Computing Machinery. ISBN 9781450359702.