

Documenting Deployment with Fabric: A Repository of Real-World AI Governance

Mackenzie Jorgensen^{*1}, Kendall Brogle^{*1}, Katherine M. Collins², Lujain Ibrahim³, Arina Shah⁴, Petra Ivanovic⁵, Noah Broestl^{2,6}, Gabriel Piles⁷, Paul Dongha⁸, Hatim Abdulhussein^{9,10}, Adrian Weller^{1,2}, Jillian Powers¹¹, Umang Bhatt^{1,2,4†}

¹The Alan Turing Institute

²University of Cambridge

³University of Oxford

⁴New York University

⁵Trustwise

⁶Boston Consulting Group

⁷HURIDOCs

⁸NatWest

⁹University of Surrey

¹⁰Health Innovation Kent Surrey Sussex

¹¹Slalom

Abstract

Artificial intelligence (AI) is increasingly integrated into society, from financial services and traffic management to creative writing. Academic literature on the deployment of AI has mostly focused on the risks and harms that result from the use of AI. We introduce Fabric, a publicly available repository of deployed AI use cases to outline their governance mechanisms. Through semi-structured interviews with practitioners, we collect an initial set of 20 AI use cases. In addition, we co-design diagrams of the AI workflow with the practitioners. We discuss the oversight mechanisms and guardrails used in practice to safeguard AI use. The Fabric repository includes visual diagrams of AI use cases and descriptions of the deployed systems. Using the repository, we surface gaps in governance and find common patterns in human oversight of deployed AI systems. We intend for Fabric to serve as an extendable, evolving tool for researchers to study the effectiveness of AI governance.

1 Introduction

Artificial intelligence (AI) is now deployed at scale, stretching from wayfinding through cities to determining creditworthiness to recommending music (Gartner, Inc. 2019; Faheem 2021; Grand View Research 2024; Mokoena and Obagbuwa 2025). Much academic literature to date has rightly focused on either advancing system capabilities or documenting risks and failures associated with system deployments. Multiple public databases, like the Responsible AI Collective’s AI Incident Database (McGregor 2021), the OECD’s AI Incidents and Hazards Monitor (Organisation for Economic Co-operation and Development 2025b), and

MIT’s AI Risk Repository (Slattery et al. 2024), highlight harms of AI system failures.

While these repositories capture AI risks, there is a dearth of knowledge on how AI systems are deployed and what governance mechanisms make such deployment safe and possible (AI.G.L. Blog 2025). To address which governance mechanisms make AI deployment less risky, the United States’ National Institute of Standards and Technology (NIST) released their AI Risk Management Framework (AI 2024). Complementary efforts include the OECD’s catalog of tools for trustworthy AI (Organisation for Economic Co-operation and Development 2025a) and the UK government’s portfolio of AI assurance techniques (Department for Science, Innovation and Technology 2023). However, the efficacy of these guidelines is unclear so far in practice.

AI systems are often quietly integrated into workforce technologies, public services, and consumer products without a shared understanding of how and when they work (John, Holmström Olsson, and Bosch 2021; Meng 2025). In this paper, we begin to fill gaps in understanding how AI systems are used in practice and what governance surrounds their use. More specifically, we ask: how is governance imposed on deployed AI systems? To address this question, we release Fabric, a public repository of AI system deployments. The deployed use cases are categorized by their sector, oversight patterns, and other meta-data.

To gather AI use cases, we conduct semi-structured interviews with AI practitioners, where we understand their AI workflows and how governance is applied. Crucially, we co-design a diagram of the underlying AI workflow with each practitioner. We present 20 AI use cases in our repository. For each use case, we describe the AI system, discuss the governance around its use, and provide the co-designed diagram. We analyze the repository to find commonalities across use cases and devise governance patterns, which we

^{*}These authors contributed equally.

[†]Correspondence to: {ncb49,usb20}@cam.ac.uk

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

hope can guide future deployment of AI. We cast a wide net in our definition of “AI systems,” ranging from large language models (LLMs) to shallow neural networks to rule-based expert systems. We focus on tactics to operationalize AI governance, the control and oversight of deployed AI systems. Instead of theorizing which principles dictate governance at scale, we capture instances of AI governance from real-world use cases.

We aim to offer academics, industry professionals, and policymakers a landscape view of how AI systems can be governed by releasing the governance patterns from our use cases. Our initial collection of them is only the beginning; we envision the repository will grow with broader community engagement. This paper proceeds as follows: we outline related work in Section 2; we describe our methodology in Section 3; we provide an overview of the Fabric repository and define the governance patterns that arise across the Fabric use cases in Section 4; we discuss the governance considerations from the use cases in Section 5; lastly, we conclude in Section 6 with our hopes for future iterations of work.

2 Related Work

AI is increasingly deployed across a wide range of decision-making contexts, from high stakes domains (e.g., healthcare (Esteva et al. 2017) and immigration (Booth 2024)) to lower-risk applications (e.g., personalized writing assistance (Hwang et al. 2023))—even potentially forming the base of new kinds of collaborative cognition in human-AI thought partnerships (Collins et al. 2024b). While AI systems offer benefits in terms of efficiency and scalability, they also introduce harmful and discriminatory risks. Thus, we need ways to govern these AI systems to mitigate for those harms and risks.

2.1 Proposals for AI Governance

AI governance can be understood as the combination of institutional, technical, and procedural mechanisms that ensure the legal, ethical, and safe development and deployment of AI systems (Jobin, Ienca, and Vayena 2019; Schneider et al. 2023), which may impact not just individuals but collective networks of people (Brinkmann et al. 2023; Collins, Bhatt, and Sucholutsky 2025). In response to risks (e.g., from biased AI systems (Mehrabi et al. 2021)), the literature proposes a range of governance mechanisms that we non-exhaustively outline.

Proposals span both legal-institutional structures and procedural mechanisms that shape how AI is used after deployment. Legal frameworks, such as the European Union’s (EU) AI Act and the EU General Data Protection Regulation (GDPR) mandate different forms of human oversight and transparency (European Parliament and Council of the European Union 2016; Edwards 2021; European Parliament and Council of the European Union 2024; AI 2024). Within the United States, NIST has created an AI risk management framework that provides recommendations but is not legally enforceable (Dotan et al. 2024). Beyond legal structures, the literature outlines several normative governance mechanisms that can be expressed post hoc, especially

when the organization has procured third party AI systems. These include human-in-the-loop review (Green 2022), escalation protocols (Upmann 2024), logging requirements (Cheong 2024), oversight by default (Centre for Information Policy Leadership 2024), post-hoc monitoring (Mökander and Floridi 2023), and workflow-based gate-keeping (Kaye, Dixon, and Gellman 2023). Unfortunately, such mechanisms are often described abstractly with little guidance as to how they can be implemented in practice.

Complementing these mechanisms, Shao et al (Shao et al. 2025) introduces a Human Agency Scale (HAS) that provides a shared, empirical language for calibrating the degree of human involvement with AI agents. Their framework reveals mismatches between worker preferences and expert assessments, mapping tasks into actionable zones of automation and augmentation. These insights can directly inform the design of oversight and escalation mechanisms so that governance aligns more closely with human preferences.

2.2 Addressing Gaps in the Literature

While existing initiatives, such as the Responsible AI Collective’s AI Incident Database, the OECD’s AI Incidents and Hazards Monitor, and MIT’s AI Risks Repository, document failures and risks associated with AI system deployment (McGregor 2021; Organisation for Economic Co-operation and Development 2025b; Slattery et al. 2025), these resources are inherently reactive. They flag where governance has failed but do not provide insight into how governance is structured in operational settings. Similarly, documentation methods, like model cards (Mitchell et al. 2019) and data sheets (Holland et al. 2020; Gebru et al. 2021; Frieder et al. 2024) promote transparency but stop short of capturing how oversight, escalation, or accountability is embedded into AI workflows.

Normative frameworks propose governance mechanisms, such as selective abstention (Chow 1970; Cortes et al. 2017), confidence disclosures (Shortliffe and Buchanan 1975; Doshi-Velez and Kim 2017; Le et al. 2023), and transparency prompts (Wachter, Mittelstadt, and Floridi 2017; Ananny and Crawford 2018; Kaminski 2021): all of which are grounded in responsible AI principles. While these proposals offer valuable guidance, there is limited empirical research showing how such mechanisms are actually implemented for real-world AI systems, or how they are practically integrated into institutional oversight structures. As a result, a critical gap remains, examining how AI governance is enacted in real-world AI systems.

This paper addresses this gap by documenting governance mechanisms across a set of deployed AI systems. By co-designing AI workflow diagrams with practitioners, we visualize how oversight, institutional policies, and accountability checkpoints operate in real-world AI use. The Fabric repository offers glimpses into what governance looks like in practice, informing both researchers and practitioners on how to bridge the divide between principles and implementation.

2.3 Repository Initiatives

The need for repositories of AI governance practices is a widely recognized necessity by governments, researchers

and industry. In fact, U.S. federal initiatives, including the National AI Initiative Act, recent executive actions, and the Department of Defense’s Task Force LIMA, have called for public registries of AI uses cases and governance practices to support responsible deployment (Department of Defense, Chief Digital and Artificial Intelligence Office 2024; Executive Office of the President 2020; McKernon et al. 2024; National Archives and Records Administration 2025; Office of the Federal Chief Information Officer 2025). Globally, repositories, such as the OECD AI Policy Observatory (Organisation for Economic Co-operation and Development 2024), UNESCO’s Ethics of AI platform (United Nations Educational, Scientific and Cultural Organization 2024), and Singapore’s Model AI Governance Framework (Infocomm Media Development Authority (IMDA) of Singapore 2024), offer centralized collections of AI-related laws, frameworks, and case studies. Industry and academic efforts have followed suit, including IBM’s AI Governance Fact-sheets (IBM Research 2024) and the University of Technology Sydney’s AI Governance Lighthouse Series (World Economic Forum 2024).

While valuable, these repositories are typically limited to policy-level summaries, template-driven documentation, or high-level descriptive accounts. They often do not examine how governance can be operationalized in existing AI workflows via human oversight, fallback mechanisms, or institutional constraints. The AI Governance Library (AI.G.L. Blog 2025) complements these repositories by offering a practitioner focused, blog style, archive of tools intended to provide professionals with practical strategies for deploying AI systems. Fabric differentiates itself from these repositories and addresses an unmet need by documenting governance structures as they are enacted. Rather than focusing solely on compliance or documentation, Fabric captures the workflow of deployment through co-designed diagrams and structured summaries of oversight mechanisms. In doing so, Fabric builds upon previous initiatives while providing a unique workflow-level, empirical, and cross-sectoral view of AI governance in action.

3 Methodology

To address the lack of empirical understanding of how governance is implemented in practice, we set out to document real-world AI deployment with Fabric. We focus on how oversight is embedded within workflows and how AI systems interact with existing infrastructure and operations at organizations. To accomplish this, we use qualitative, semi-structured interviews (Gubrium and Holstein 2002) and collaborative co-design (Poinet et al. 2020) to depict the elements that shape system governance. Co-design, also dubbed prototyping, is a well-accepted methodology for creating artifacts alongside a participant (Buchenau and Suri 2000; Madden et al. 2014; Camburn et al. 2017): here, we work with an AI practitioner to create AI workflow diagrams similar to that shown in Figure 1. We obtained ethics approval for this study from The Alan Turing Institute’s Institutional Ethics Review Board.

To gather an initial set of use cases, over the course of five months, we circulated more than 80 invitations to partici-

pate, not including those sent through institutional mailing lists. We conducted virtual interviews with practitioners, individuals directly involved in the development, deployment, or maintenance of AI systems within an organization. Each semi-structured interview lasted around 45 minutes. We include our interview questions in Appendix A. During the interview, we collaborated with the practitioner to create an AI workflow diagram. The practitioners described their workflow, while we mapped out an initial diagram to review with the practitioner. Each interview had at least one interviewer present (most had two): one person asking questions and the other iterating on the diagram based on practitioner feedback. Real-time collaboration let us clarify doubts immediately and incorporate the practitioner’s feedback iteratively to ensure the diagram represented the workflow well. After the interview, we refined the diagram and wrote a use case summary to share with the practitioner for approval. Once we collected multiple use cases, we sought governance patterns that emerge across workflows.

3.1 Sampling and Interview Subjects

Participants were recruited through personal networks, research partnerships, and Slack channels. We expanded our pool of interviewees through snowball sampling (Biernacki and Waldorf 1981; Noy 2008; Parker, Scott, and Geddes 2019). We adopted a broad definition of “AI system” to capture a wide audience. Before each interview, participants received a detailed study overview and signed a consent form that outlined our ethics considerations and privacy measures. Given our sample size, we used pseudoanonymization during data collection and storage to preserve anonymity. We note that interviews were optionally recorded and later deleted, after transcriptions of them were pseudoanonymized and saved.

In this paper and in the initial Fabric repository, we do not disclose the names of organizations from which the use cases came, unless they specifically asked and provided us with consent to share this information. Each use case is categorized by organizational sector (e.g., private, public, and civil society) and domain, and all identifying information is stored securely and only accessible to the research team.

3.2 Diagram Design

We developed a “diagram design” approach to visually represent both technical workflows that include AI systems and the governance mechanisms at play. While we initially drew inspiration from software modeling frameworks, such as Unified Modeling Language (UML) diagrams (Jäger, Schleicher, and Westfechtel 1999; Jacobson and Booch 2021), we ultimately chose to use decision flowchart design to incorporate elements such as decision points, flow lines, and feedback loops (Yonyx 2024; ProjectManager 2025). Decision flowchart design represents the interactive nature of AI systems and governance.

As we interviewed practitioners, we updated the elements included in the diagrams because we needed to account for more complex workflows. We could not represent all elements of the workflows, so we prioritized certain common elements and included more fine-grained details in the use

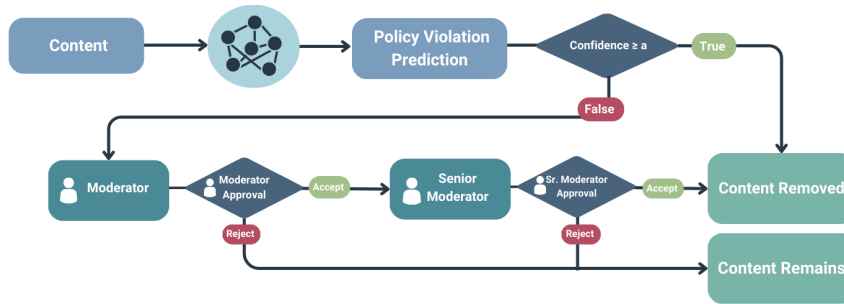


Figure 1: An example AI workflow for a hypothetical content moderation system deployed by a media platform to evaluate user content.

case summaries. We completed the diagram drafting on the Canva platform. A comprehensive key of the elements used in our diagrams is included in Figure 1 in Appendix B.

To ground our methodology in a concrete example, we walk through a representative use case for a content moderation system used by a hypothetical social media platform to identify content policy violations. In Figure 1, the content is the input for an AI system. The AI system then evaluates if the content violates the platform’s policy. In addition to a prediction of content violation, the AI system also outputs a confidence score for the prediction. The next step in the workflow is a decision point which occurs as a safety check on the system’s confidence score. A decision point is a place in the workflow that determines the next course of action. In this case, if the confidence score is above a predetermined threshold, then the content is automatically removed without any human intervention; thus, the AI system could act autonomously. However, if the score falls below the threshold, then the content is routed to a human moderator for review. If the moderator does not believe the content violates their policies so it is deemed safe, then the content remains on the platform. If the moderator believes the post violates their policy, then an additional moderator has to check this violating content before it is removed. The “final output” of this workflow is either the content is removed or the content remains, both of which is shown the diagram. We use these components to generate diagrams for the initial repository.

3.3 Data Analysis

To analyze what patterns of governance arose, our research team reviewed interview notes and AI workflow diagram summaries, and compared diagrams across use cases. As use case summaries and diagrams were built, we iteratively reviewed our findings to update the patterns of governance. In subsequent sections, we catalog the patterns found and report how many AI use cases fall under each pattern.

4 Fabric: The Initial Inventory

The initial release of our Fabric repository is made up of 20 real-world use cases for AI, collected through semi-structured interviews with practitioners across a variety of

sectors and types of organizations. Each use case centers around an AI system that is in active use within an organization. Repository entries include: the sector, domain, an AI workflow diagram that shows the operation AI is involved in, a brief description of the task, the intention behind deploying the system, its risks, patterns of governance, and an explanation of the AI workflow diagram. Collectively, these use cases are intended to ground abstract theories on AI oversight in empirical reality, documenting how governance is enacted in practice. In Appendix B, we show all 20 use cases from the initial Fabric repository, including their diagrams and descriptions. We also include summary statistics of the use cases in Table 1 in Appendix B.

4.1 Patterns of AI Governance

From our initial repository of 20 use cases, we expected three main types of AI systems to emerge with varying degrees of human oversight: (1) fully-automated systems where there is no human oversight, (2) human in the loop systems where the human can accept or modify AI outputs, and (3) human in the loop systems where the human can only accept or reject AI outputs, not change them. However, on review of the AI use cases, our analysis resulted in four different levels of governance which we will describe in detail below. We categorize these levels of governance on top of AI systems as types of *human oversight* governance. To determine an AI system’s human oversight level, we ask: can a human affect the final output of the system? If so, what actions can they take to determine the final output?

As we reviewed the mechanisms of governance across use cases, another type of governance emerged; we name this form of governance: institutional oversight. Practitioners shared that institutional oversight can happen internally within AI systems’ processes or externally. We define (see Table 1) and provide examples of the different levels of human oversight and institutional oversight in our initial Fabric repository. We include Figure 2 which groups each use case by their oversight level and the strictest form of institutional oversight present in the system. For a table of the use cases and their oversight patterns, see Table 2 in Appendix B. While the specific governance types and levels we present

Level	Institutional Oversight	Definition
Low	Ad-Hoc Practice	A mechanism subjectively determined by a developer or user.
	Organization Best Practice	A practice determined by the organization or teams inside of it.
	Organization Policy	A policy requiring certain practices to be followed.
	Industry Standard	A standard from industry best practice or an official standards body.
High	Regulation	The rules and processes necessary to enable a law.
Human Oversight		
Low	Autonomous AI	AI output is the final output.
	Conditionally Autonomous AI	AI output could be the final output; otherwise, output is up to the human.
	Human-Approved AI	AI output must be approved by a human as the final output.
High	Human-Led with AI-Assistance	AI output could be used by a human in a larger final output.

Table 1: Definitions of institutional oversight and human oversight levels from low to high levels of governance strictness.

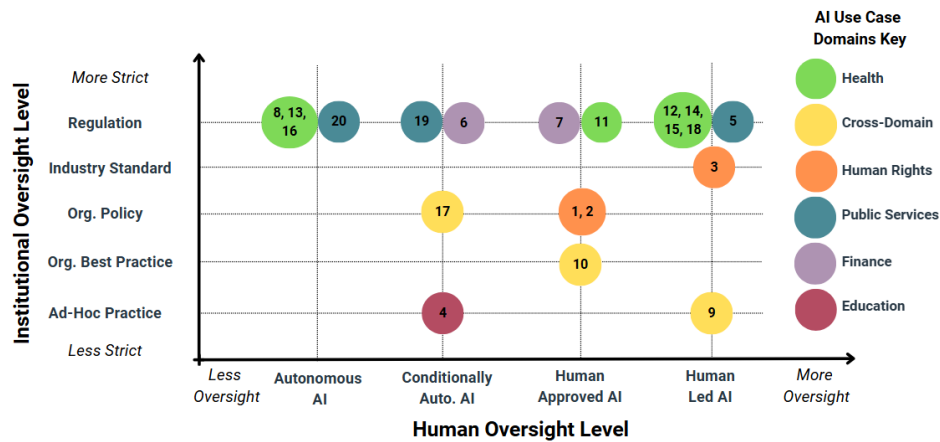


Figure 2: The oversight levels across all of the use cases organized by their domains. The conditionally autonomous AI level is included as “Conditionally Auto. AI” and the human-led with AI-assistance level is written as “Human-Led AI.” We plot the strictest form of institutional oversight for each use case. The numbers match the use case numbers in Appendix B.

are not entirely novel, our contribution is in clear labels, consistent structure, and most importantly in documenting the application of governance in deployed AI systems. When referring to the use cases in the following text, we include their number in the repository to find them in Appendix B.

4.2 Institutional Oversight Levels

Within the initial Fabric repository, we found multiple governance mechanisms that did not fall under the human oversight levels we uncovered. These institutional oversight mechanisms included a broad range of interventions, such as metric threshold checks, internal policies, industry standards, and regulatory compliance. We define the different levels of institutional oversight from less strict governance (e.g., ad-hoc practice) to more strict governance (e.g., regulation) in Table 1. We now outline examples of institutional oversight mechanisms from the repository for each governance level. Use cases can include multiple levels of institutional oversight in a single AI system workflow.

Ad-Hoc Practice This level of institutional oversight is the least strict, as individuals or teams decide governance

“off-the-cuff,” often uninformed by internal guidance or external pressures. Within the repository, two use cases included only ad-hoc practices (the personalized feedback assessor No. 4 and Writing Assistant No. 9) and no stricter levels of institutional oversight were applied.

- Personalized Feedback Assessor No. 4: This AI system generates feedback for students based on their online quiz performance. Before sending the feedback to the students, certain safety conditions must be met, including checking if the answer is scientifically correct or if harmful language is used. This safety check is an ad-hoc oversight practice since it is up to the developer to decide what those checks are and what thresholds they should be. If a safety check falls below a certain threshold, the feedback is sent to the student’s teacher and the AI system regenerates new feedback.
- Writing Assistant No. 9: The organization that provided the personalized writing assistant use case used a third party to procure their AI system. They chose to apply sensitivity filters from the third party to their AI system. As an ad-hoc oversight measure, the organization opted

to implement the least restrictive sensitivity filter setting.

Organization Best Practice Many practitioners considered their organizational best practices to be an adequate form of institutional oversight. Some of these practices include confidence thresholds set for safety checks (e.g., Insurance Claims Classifier No. 6) and prioritization of certain patients if labeled as high-risk from an AI system (e.g., the CT-Scan Risk Detector No. 14).

- Insurance Claims Classifier No. 6: In the insurance claim workflow, there is a denial confidence threshold check, determined by the insurance organization. This check determines if the case should be considered by a claims analyst or if the AI is allowed to make the final decision.
- CT Scan Risk-Detector No. 14: For this example, the AI system flags anomalies of the scan and provides a priority score, after which radiologists see the high-risk scans first. This lets radiologists process next steps for patients who need it most.

Organization Policy Organizations often institute internal policies specific to their sector and domain (e.g., Mental Health Triage Tool No. 13 and OriginTrail Decentralized Knowledge Graph No. 17). If these policies are not followed, then sometimes employees could face repercussions.

- Mental Health Triage Tool No. 13: In the mental health triage tool, we found several examples of governance via internal policy. For a brief insight into the tool, it provides relevant forms for a user to fill out and organizes a preliminary patient report for healthcare providers. The organization policy outlines that a lived-experience digital group must approve a system before it is deployed. A project board, clinical/operational design authority, and technical designer all must sign off before deployment. By including checkpoints, risks are ideally mitigated before deploying systems.
- OriginTrail Decentralized Knowledge Graph No. 17: This use case provides knowledge graphs to the public. Users can build their own knowledge graphs. The organization has an open-source policy for their code so that anyone can build knowledge graphs and gain insights from them. Their open-source policy can result in the community flagging safety and risks of their software, supporting governance.

Industry Standard These are standards that could come from industry best practice or a national or international standards body (e.g., International Organization for Standards). Most examples of industry standards came from the healthcare use cases, including medical device standards (e.g., Physio-Risk Reporter No. 8, Mental Health Triage Tool No. 13, and Carer-AI Kit Risk Assessor No. 16). However, other use cases did cover other industry standards; accessibility standards were a priority for a Local Government Chatbot No. 5.

- Local Government Chatbot No. 5: A local government organization built a chatbot for residents to find information about less risky topics on their website. The organization was aware that they needed to make their chatbot

accessible for users with visual and other impairments. Given this priority, they followed accessibility standards for their chatbot.

- Mental Health Triage Tool No. 13: This use case provides relevant forms for a user to fill out and organizes a preliminary patient report for healthcare providers. Some of the relevant standards implemented include: safe deployment of healthcare technology with designated safety officers, and a health and equality impact assessment.

Regulation This level of institutional oversight is the strictest because regulations are enforceable by law, and failure to comply can lead to legal consequences including sanctions, litigation, or regulatory action

Our use cases captured a broad range of regulations from domains like finance (e.g., Credit Lending Classifier No. 7), healthcare (e.g., Physio-Risk Reporter No. 8), and public services (e.g., Local Government Chatbot No. 5). Almost all of the public sector organization use cases (e.g., Carer-AI Kit Risk Assessor No. 16 and Mental Health Triage Tool No. 13) that process or store sensitive data shared their compliance with data protection regulation as key institutional oversight for their AI systems. We note that how organizations abide by regulation is usually organization dependent.

- Credit Lending Classifier No. 6: The AI system in this use case had to comply with data protection regulation since it accesses personal data about credit applicants. Among other practices to comply, this organization had to complete data protection impact assessments.
- Chest X-Ray Abnormality Detector No. 12: This use case involves radiation so must follow radiation regulation. Radiation regulation ensures the safety of patients and those running the scans so is a crucial form of institutional oversight of AI systems.

4.3 Human Oversight Levels

We discuss human oversight levels that we observed in the initial Fabric repository. Upon conducting our interviews, we noted that our expected “human in the loop” AI system types were quite similar and, at times, overlapping. However, we found a key distinction between these two types of AI systems and the human oversight applied—in some instances, the human could update the data given to the AI system to either retrain it or rerun with the same data and system, hoping for different outputs. To account for the nuances in the repository, we redefine the “human in the loop” categories into the following human oversight levels: (1) human-approved AI and (2) human-led with AI-assistance.

We provide specific examples of the human oversight levels observed in Fabric. For each level, we show an example diagram, describe the key abstract features of the oversight level, and discuss at least one of the observed use cases representative of that level.

Autonomous AI Below, we discuss the autonomous AI governance pattern, see Figure 3 for a visual. Features that fall under this human oversight level include:

- The AI output is the final output.



Figure 3: The autonomous AI diagram which is the least strict level of human oversight.

- No human oversight can change the final output.

From the initial Fabric repository, we categorize four use cases as autonomous AI: Physio-Risk Reporter No. 8, Mental Health Triage Tool No. 13, Carer-AI Kit Risk Assessor No. 16, and Image Blurring Tool No. 20. Every autonomous AI use case within Fabric also includes some form of institutional oversight.

- Physio-Risk Reporter No. 8: In this use case, an AI system provides a user (over 65 years old) with a risk score and a report regarding their physio-related status. The user completes prescribed exercises with a physiotherapist (institutional oversight example) while a video is recorded. This video is fed into an AI system, which includes two distinct models before the final output, the risk report, is presented to the user. There is human oversight present for the first AI model but not for the second. A physiotherapist reviews the first models output: a skeletal animation of the user’s movements as they complete their exercises. If the skeletal video quality is not adequate, then the physiotherapist can request that the user re-record. We note that there is no governance on the second AI model, the risk report generator. The AI output is the final output aligning this use case with the autonomous AI level.

Conditionally Autonomous AI For the conditionally autonomous AI oversight pattern, see the top diagram in Figure 4. The use cases that fall under this human oversight level share the following features:

- The AI system can act autonomously.
- If the AI system cannot act autonomously given certain constraints, then the human must make the ultimate decision for the final output.

We observe four use cases in Fabric that are categorized under this oversight level: Personalized Feedback Assessor No. 4, Insurance Claims Classifier No. 6, OriginTrail Decentralized Knowledge Graph No. 17, and Call Center Virtual Assistant No. 19. The determining factor for who makes the final decision, the AI or a human, is a safety check which could include a metric or confidence check of the AI output. We note that if the human must participate in the process, we see human oversight like human-approved AI.

- Personalized Feedback Assessor No. 4: The AI system provides students with feedback on their quiz answers. The feedback is generated and given to the student automatically, so the AI system acts autonomously. The AI acts autonomously unless, a teacher has chosen that they want to check the AI feedback before giving it to students. In this instance, we then see a type of human-led

with AI-assistance hybrid type of governance, where the teachers accept the feedback or they can edit the AI output before its sent to the student.

- Insurance Claims Classifier No. 6: Unlike the personalized assessment feedback tool where the teacher determines if they want the AI to run autonomously, the insurance claim use case has *two* decision points to determine whether or not the AI system runs autonomously or if human decision-makers need to review the claim. If the claim is simple and the AI system has a high denial confidence (institutional oversight example), then the claim is denied and the AI system has acted autonomously. If the claim is not simple or if the AI system’s denial confidence is below a certain threshold, then a claims analyst has to decide if its approved or denied. Whether its an AI or a claims analyst, a customer could appeal the claim as a matter of institutional oversight but this is after the final output.

Human-Approved AI For the human-approved AI governance pattern, see the middle diagram in Figure 4. One of the two human in the loop scenarios must be true for the use case to fall under this human oversight level:

- A human can accept the AI output as the final output or reject it (and not have a final output).
- If the human is not satisfied with the AI output, then they can modify the input data to potentially receive a different and better output.

We have five use cases that fall under this level of governance: Metadata Extractor No. 1, PDF Segmentator No. 2, Credit Lending Classifier No. 7, Hyperparameter Optimizer for LLM and RAG Systems No. 10, and Smart Clinical Triage Tool No. 11.

- Metadata Extractor No. 1: The purpose of the AI system is to build a model that extracts user-specified metadata from samples. The user then decides if they are satisfied with the model’s predictions and if so, they can use the trained model for their tasks, if not, they can delete the model or update their samples and train a new model. This use case is particularly interesting because it is one of the few uses cases we collected where a trained model is the *final output*.
- Credit Lending Classifier No. 7: The AI system generates a recommendation for credit officers to review before they make the final decision and there are different variables (e.g., loan size) that determine how many credit officers (institutional oversight example) must review the loan application before the final outcome is determined. Ultimately, the credit officers have the final say if they agree or reject the AI system’s recommendation.

Human-Led with AI-Assistance For this oversight level see the bottom diagram in Figure 4. In this level, the human always makes the final decision with AI-Assistance. The features that fall under this human oversight level include:

- The AI system output can be used by the human, if they are satisfied and agree with it.

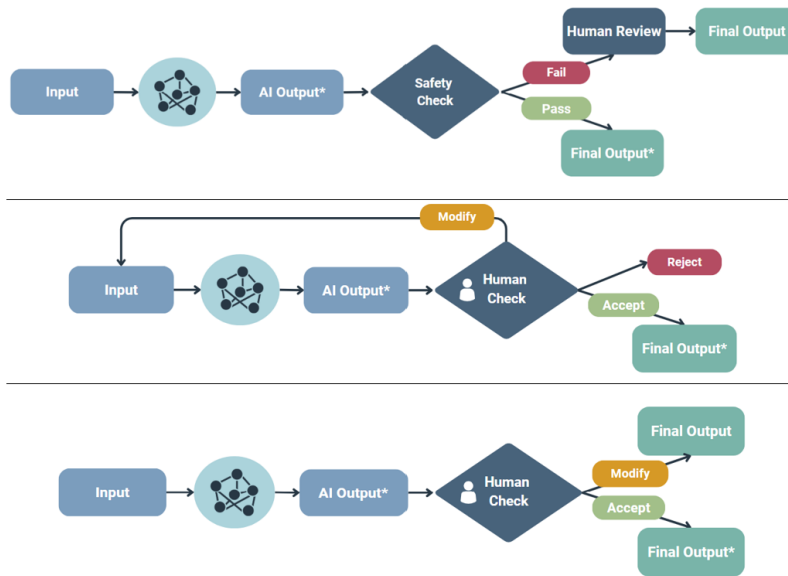


Figure 4: We include three of the human oversight levels in visual form. The top diagram is a conditionally autonomous AI system. The middle diagram is a human-approved AI system. The bottom diagram is a human-led with AI-assistance system. We denote joint human-AI approved final outputs with a *.

- The final output of the system is more than the AI system output and involves more work by the human.

We have seven use cases that fall under this level of oversight: Family Court Support Chatbot No. 3, Local Government Chatbot No. 5, Writing Assistant No. 9, Chest X-ray Abnormality Detector No. 12, CT Scan Risk Detector No. 14, Consultation AI Note-Taker No. 15, and Automated Imaging Protocol Selector No. 18.

At this level of human oversight, the human can decide to keep working with the support of the AI (e.g., Family Court Support Chatbot No. 3) or the human can decide to use the AI output for their final decision (e.g., Chest X-Ray Abnormality Detector No. 12, and CT-Scan Risk Detector No. 14).

- Family Court Support Chatbot No. 3: The human can keep chatting with the chatbot tool until they either have their answer or decide to take their query to a human for support. This use case has loops in the functionality so the human can keep requesting things of the AI system.
- CT Scan Risk Detector No. 14: The radiologist who views the initial scan and receives the secondary capture from the AI system makes uses of their own knowledge and the AI output if they agree with it in their final report; hence the AI assistance is there if deemed useful by the human.

5 Discussion

We outline the main takeaways from the curation and analysis of the initial Fabric repository. We also discuss the value and limitations of our particular methodology.

5.1 Patterns of Oversight

The levels of human oversight and institutional oversight discussed above, together, make a “Pattern of Governance” and more specifically, a “Pattern of Oversight.” Through our study, we saw these “Patterns of Oversight” which are the decisions of the two levels of oversight paired together. We argue that these patterns should be included in model cards and documentation for transparency purposes. Currently, much of the thinking behind governance happens post-development of the AI system but pre-deployment. We encourage the community to discuss these patterns pre-development and to revisit them continuously throughout the AI life cycle. In Appendix C, we include a questionnaire to support individuals or teams to check what level of human oversight their AI use case falls under; in addition, we provide some questions to support discussions as to what institutional oversight might be necessary given different risks.

We outline common patterns of oversight that we see in our repository below. We note that many of our use cases are considered high-risk in line with the EU AI Act’s framework; high-risk systems are those used within contexts where failure can significantly harm health, safety, or fundamental rights (European Parliament and Council of the European Union 2024). It is therefore unsurprising that many of them followed the strictest level of institutional oversight.

- The most common human oversight pattern is “human-led with AI assistance,” so a decision-maker can modify or accept an AI’s output as a final output. The most frequent level of institutional oversight was regulation for these use cases, which was also the most common institutional oversight across all use cases.
- All AI use cases that fall into the autonomous AI over-

sight level abide by the strictest form of institutional oversight, regulation.

- All conditionally autonomous AI use cases, except for the more high-risk financial use case, follow the least strict level of institutional oversight: “ad-hoc practice.”

5.2 Institutional Oversight Considerations

The majority of the practitioners that we interviewed for our repository could not point to best practice AI governance guidelines. This lack of clarity on institutional oversight could reflect either an absence of sufficient governance, or a design flaw in the governance process that prevents actors within the organization from being aware. In both cases, the use of AI is not effectively guided or constrained. However, practitioners could point to regulation or domain-specific industry standards that they followed. The majority of use cases included at least some form of regulatory compliance in their AI systems. This was especially clear for government use cases. In less risky domains, ad-hoc practices were common while stricter levels of governance were often absent, suggesting that practitioners may have struggled to integrate emerging AI governance and safety measures alongside industry requirements.

We found that most of the institutional oversight was decided in-house to comply with regulations, industry standards, organizational policies, or general “best practices.” Some practitioners acknowledged that while they did not have specific AI governance in place when the AI systems were built, they are working on or have governance frameworks in place internally now. This was the case for the Local Government Chatbot No. 5 use case. This finding highlights how the practitioners across industries (at least in our recruited pool) are aware of the risks posed by using AI systems in their processes and are prioritizing best practices or policy frameworks internally for current and future AI use once the technical envelopment of their use cases are production/user ready.

Many public and private sector use cases leveraged AI systems that were not developed internally but were procured from third-parties. Those deploying third-party systems had less control of system innards. The choice to use third-party AI systems stems from the appeal of accessibility and ease of use. Perhaps more concerning, there is a lack of visibility across organizations that may make attempts to debug, escalate, and correct particular failures challenging.

Some interviewed practitioners expressed the need for more knowledge sharing of what AI governance can look like *across* industries and sectors. They were curious to see how others had implemented AI systems and included forms of governance. This interest from the practitioners themselves points to the fact that the Fabric repository resonates and will provide valuable examples of governance in action.

In practice, institutional oversight still varies widely. There are several reasons for this, including:

- The locus of control (primarily in either the Chief Tech Officer or Chief Risk Officer) is not standardized, often resulting in misaligned governance.

- International, national, and local legislation, which can differ significantly, is evolving (or in some cases is completely absent), which creates a lack of clarity and a hesitance to create heavy-weight processes that might not respond to regulatory shifts.
- There has yet to be a body of generally accepted practices that has been developed around AI, as the field is in its infancy, creating even more uncertainty in the space.

Fabric may provide support to facilitate alignment, clarity, and, in future iterations, establish best practices.

5.3 Human Oversight Considerations

As noted before, all four of the autonomous AI use cases follow the strictest level of institutional oversight (see Figure 2 above, and Table 2 in Appendix B). This is notable because these systems are subject to greater levels of regulatory requirements, suggesting they may be higher risk, yet they have the least amount of human oversight. For the conditionally autonomous AI systems (when the output is not fully autonomous), we see the human-approved AI pattern appear in all use cases. The organizations we survey recognize that there can be a combination of human oversight patterns within a system process. Most of the use cases in Fabric fall under the human-led AI assistance human oversight category. This category includes the highest level of human oversight with an AI system still “in the loop.” In this instance, the human can decide if the AI output should contribute to the final output, which requires human labor beyond the potential AI contribution. Although this potentially requires more labor from the human involved, such a governance structure can allow for a strong safeguard before integrating, if at all, the AI output into a decision (Steyvers et al. 2022; Collins et al. 2024a). We highlight that human oversight levels give rise to a set of controls that must be followed; the deciding factor of what level is necessary often depends upon the risk rating of the use case.

We note that the human could be over- or under-reliant on the AI system and result in negative outcomes. Specifically, the concept of automation bias, when users are prone to accept AI-generated outputs without adequate critical evaluation, and the harms are well-documented (Skitka, Mosier, and Burdick 1999; Gosline et al. 2024). The consequences of over-reliance on AI can also negatively impact the governance of such systems (Sauer, Kao, and Wastell 2012). For instance, if the human over-relies on the system and does not adequately check the output for errors, they could include a mistake in the final output, like leading to a misdiagnosis of a patient (Logg, Minson, and Moore 2019). On the other hand, the AI system could provide a useful output like a critical anomaly in a chest X-ray that the doctor had missed; the doctor could completely overlook the AI’s output because they under-rely on the system, resulting in a less robust final report being produced (Bonaccio and Dalal 2006). Fabric use cases that fall under the human-led AI assistance and human-approved AI levels are common in the literature when studying what sorts of human oversight or technical guardrails could be useful to apply for less bias or uncertain outcomes (Hou and Jung 2021; Lai et al. 2023).

5.4 Values and Limitations of Our Study

While we cannot conclude anything about the effectiveness of human oversight and the downstream impacts from an analysis of Fabric alone, we argue that focusing on oversight is crucial to understand how governance looks like in practice. Future research must study the effectiveness of each mechanism in our repository. We hope to then link these results in our project website¹ later to provide more insights for practitioners.

An inherent value in our approach is that we used qualitative methods for data collection, allowing for more unique insights to emerge. For instance, some practitioners did not initially consider how their AI system included multiple ML models until midway into the interview (e.g., Physio-Risk Reporter No. 8 and PDF Segmentator No. 2). Some practitioners seemed to hold an abstract view of AI in organizational processes and did not consider the fine-grained specification of the AI system as a whole. This points to differences in the practitioner perception of AI versus its actual role within the workflows, and the value of open-ended interviewing and diagram development for the discovery and nuance of oversight.

We also highlight that co-designing use case diagrams in the interviews was an effective practice to open up the discussion about the overall system. We found that the diagram discussions led to the best understanding of the human oversight and institutional oversight patterns. The more interviews we conducted, the more attuned we were to know what kind of questions answered and data we needed to build the repository and analyze the oversight patterns.

Since we used our personal and professional networks to recruit practitioners as participants, we encounter selection bias in our recruitment; although, we attempted to mitigate some of this concern with snowball sampling. In future iterations of Fabric, we hope to recruit a more representative sample. We see the double-edged sword of our semi-structured interviews; while such personal conversations are highly information-rich, they are difficult to scale.

In next iterations of Fabric, we plan to update our interview questions to include more specific questions about third-party AI systems and different aspects of institutional oversight. The earlier interviews we ran were less focused on institutional oversight, as this type of governance became more apparent later on in our interviewing. To mitigate this discrepancy, we followed up with practitioners and received varying levels of detail. We also note that the oversight mechanisms are often domain-specific; as such, domain-specific findings may become clearer in future iterations of Fabric. As AI systems themselves adapt, we anticipate continually revisiting and adapting Fabric, particularly as new levels or patterns of oversight emerge. Fabric could also be married with other approaches for tracking changes in AI systems like FeedbackLogs (Barker et al. 2023).

6 Conclusion

In this paper, we release Fabric, a public repository of AI use cases and corresponding governance. We collect examples

¹<https://fabric-repository.github.io>

of deployed AI systems from practitioners through semi-structured interviews. In these interviews, we co-design diagrams of a practitioner’s AI system and corresponding governance. Through these interviews, we aimed to understand oversight is implemented in practice and where these mechanisms lie within AI workflows. This knowledge fills a crucial gap in the community’s knowledge base. We discuss two types of oversight that emerge from the use cases that we collate as part of Fabric: human oversight and institutional oversight. Human oversight relates to how much oversight of the AI system occurs immediately before a final output is determined. Institutional oversight is linked to interventions that occur along the system workflow that add layers of safety to ensure the reliability and trustworthiness of the system.

From these interviews, we curate and release Fabric, a deployed AI use case repository, that leverages easy-to-interpret diagrams to characterize how AI systems are governed by different levels of oversight. We release 20 initial use cases in the repository. In the next iteration of Fabric, we will add functionality to the repository for other practitioners to add their own use cases. This will allow Fabric to grow. As we collect more use cases, we garner more insights about AI use at scale in the real-world. We hope to see other patterns of oversight arise in future iterations of Fabric. Crucially, we will begin assessing the efficacy of different governance mechanisms on target outcomes. By growing our repository, we can foster broader societal and academic awareness of the ever-evolving use of AI in the *fabric* of our everyday life.

Acknowledgments

The work of MJ, KB, UB, and AW is supported in part by ELSA - European Lighthouse on Secure and Safe AI funded by the European Union under grant agreement No. 101070617. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the European Commission can be held responsible for them. AW also acknowledges support from a Turing AI Fellowship under EPSRC grant EP/V025279/1, The Alan Turing Institute, and the Leverhulme Trust via CFI. KMC acknowledges support from the Cambridge Trust and King’s College Cambridge. We would like to thank Diletta Huyskes for her support in the early stages of the project and Sophia Worth for her feedback on the paper. We also greatly appreciate the contributions that Michael Convey, Runcong Zhao, Jurij Skornik, Michelle James, Sameer Gangoli, Koen Cobbaert, Jenny Partridge, and Enrico Ferraris made for Fabric.

Ethical Considerations Statement

While we foresee ample positive benefits of Fabric based on the practitioners we interviewed, we understand that there are privacy risks, some of which we aimed to mitigate by keeping the organization and practitioner name anonymous (unless the practitioner specifically gave their consent and requested for their organization to be de-anonymized). We

did make the risk clear to the practitioners that they could be re-identified from the repository, even though we took necessary precautions. Though we do not suspect this, practitioners could have misled us about the governance that they have in place; to some extent, regulatory compliance is key, so most organizations want to project strength and purport adherence. That said, we have no way of verifying or auditing the correctness of a practitioner's claims, but do have reviewed sign-offs for each of the use cases.

Although not a focus of our study, people affected by the AI systems we analyze will benefit from gaining an increased awareness of what guardrails and oversight are included in AI systems. There is no direct risk from our research for the people affected by AI systems in our repository; we do not include any specific instances of them or have access to any of their data.

In addition, AI practitioners (those we have not met with), by reviewing our repository, can better understand how governance can look like in practice and implement it accordingly for their own AI use cases. However, there is a risk that they over-rely on the repository instead of critically thinking about what governance mechanisms make the most sense or will be most effective for their use case.

References

- AI, N. 2024. Artificial intelligence risk management framework: Generative artificial intelligence profile.
- AI.G.L. Blog. 2025. The AI Governance Library.
- Ananny, M.; and Crawford, K. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society*, 20(3): 973–989.
- Barker, M.; Kallina, E.; Ashok, D.; Collins, K.; Casovan, A.; Weller, A.; Talwalkar, A.; Chen, V.; and Bhatt, U. 2023. FeedbackLogs: Recording and incorporating stakeholder feedback into machine learning pipelines. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–15.
- Biernacki, P.; and Waldorf, D. 1981. Snowball sampling: Problems and techniques of chain referral sampling. *Sociological methods & research*, 10(2): 141–163.
- Bonaccio, S.; and Dalal, R. S. 2006. Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, 101(2): 127–151.
- Booth, R. 2024. 'AI' tool could influence Home Office immigration decisions, critics say. *The Guardian*.
- Brinkmann, L.; Baumann, F.; Bonnefon, J.-F.; Derex, M.; Müller, T. F.; Nussberger, A.-M.; Czaplicka, A.; Acerbi, A.; Griffiths, T. L.; Henrich, J.; et al. 2023. Machine culture. *Nature Human Behaviour*, 7(11): 1855–1868.
- Buchenau, M.; and Suri, J. F. 2000. Experience prototyping. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*, 424–433.
- Camburn, B.; Viswanathan, V.; Linsey, J.; Anderson, D.; Jensen, D.; Crawford, R.; Otto, K.; and Wood, K. 2017. Design prototyping methods: state of the art in strategies, techniques, and guidelines. *Design Science*, 3: e13.
- Centre for Information Policy Leadership. 2024. Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework. Technical report, Centre for Information Policy Leadership, Brussels.
- Cheong, B. C. 2024. Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making. *Frontiers in Human Dynamics*, 6: 1421273.
- Chow, C. 1970. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1): 41–46.
- Collins, K. M.; Bhatt, U.; and Sucholutsky, I. 2025. Revisiting Rogers' Paradox in the Context of Human-AI Interaction. *arXiv preprint arXiv:2501.10476*.
- Collins, K. M.; Chen, V.; Sucholutsky, I.; Kirk, H. R.; Sadek, M.; Sargeant, H.; Talwalkar, A.; Weller, A.; and Bhatt, U. 2024a. Modulating language model experiences through frictions. *arXiv preprint arXiv:2407.12804*.
- Collins, K. M.; Sucholutsky, I.; Bhatt, U.; Chandra, K.; Wong, L.; Lee, M.; Zhang, C. E.; Zhi-Xuan, T.; Ho, M.; Mansinghka, V.; et al. 2024b. Building machines that learn and think with people. *Nature human behaviour*, 8(10): 1851–1863.
- Cortes, C.; Gonzalvo, X.; Kuznetsov, V.; Mohri, M.; and Yang, S. 2017. Adanet: Adaptive structural learning of artificial neural networks. In *International conference on machine learning*, 874–883. PMLR.
- Department for Science, Innovation and Technology. 2023. Portfolio of AI Assurance Techniques.
- Department of Defense, Chief Digital and Artificial Intelligence Office. 2024. Task Force Lima Executive Summary. Technical report.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Dotan, R.; Blili-Hamelin, B.; Madhavan, R.; Matthews, J.; and Scarpino, J. 2024. Evolving AI Risk Management: A Maturity Model based on the NIST AI Risk Management Framework. *arXiv preprint arXiv:2401.15229*.
- Edwards, L. 2021. The EU AI Act: a summary of its significance and scope. *Artificial Intelligence (the EU AI Act)*, 1: 25.
- Esteva, A.; Kuprel, B.; Novoa, R. A.; Ko, J.; Swetter, S. M.; Blau, H. M.; and Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639): 115–118.
- European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).

- Published in the Official Journal of the European Union, L 119, 27 April 2016.
- European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Published in the Official Journal of the European Union, L 259, 12 July 2024.
- Executive Office of the President. 2020. Executive Order 13960: Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government. Technical report, The White House.
- Faheem, M. A. 2021. AI-Driven Risk Assessment Models: Revolutionizing Credit Scoring and Default Prediction. *Iconic Research And Engineering Journals*, 5(3): 177–186.
- Frieder, S.; Bayer, J.; Collins, K. M.; Berner, J.; Loader, J.; Juhász, A.; Ruehle, F.; Welleck, S.; Poesia, G.; Griffiths, R.-R.; et al. 2024. Data for mathematical copilots: Better ways of presenting proofs for machine learning. *arXiv preprint arXiv:2412.15184*.
- Gartner, Inc. 2019. Gartner Survey Shows 37 Percent of Organizations Have Implemented AI in Some Form. *Gartner Press Release*.
- Geburu, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Gosline, R. R.; Zhang, Y.; Li, H.; Daugherty, P.; Chakraborty, A. D.; Roussiere, P.; and Connolly, P. 2024. Nudge Users to Catch Generative AI Errors. *MIT Sloan Management Review*, 65(4): 22–24.
- Grand View Research. 2024. Artificial Intelligence Market Size, Share & Trends Analysis Report By Solution, By Technology, By End Use, By Region, And Segment Forecasts, 2023 - 2030. *Grand View Research*.
- Green, B. 2022. The flaws of policies requiring human oversight of government algorithms. *Computer Law & Security Review*, 45: 105681.
- Gubrium, J. F.; and Holstein, J. A. 2002. *Handbook of interview research: Context and method*. Sage.
- Holland, S.; Hosny, A.; Newman, S.; Joseph, J.; and Chmielinski, K. 2020. The dataset nutrition label. *Data protection and privacy*, 12(12): 1.
- Hou, Y. T.-Y.; and Jung, M. F. 2021. Who is the expert? Reconciling algorithm aversion and algorithm appreciation in AI-supported decision making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–25.
- Hwang, W.-Y.; Nurtantyana, R.; Purba, S. W. D.; Hariyanti, U.; Indrihapsari, Y.; and Surjono, H. D. 2023. AI and recognition technologies to facilitate English as foreign language writing for supporting personalization and contextualization in authentic contexts. *Journal of Educational Computing Research*, 61(5): 1008–1035.
- IBM Research. 2024. AI FactSheets: Enabling Trust in AI. <https://www.research.ibm.com/interactive/ai-factsheets>. Accessed: 2025-05-23.
- Infocomm Media Development Authority (IMDA) of Singapore. 2024. Model AI Governance Framework. <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2024/Model-AI-Governance-Framework>. Accessed: 2025-05-23.
- Jacobson, L.; and Booch, J. R. G. 2021. The unified modeling language reference manual.
- Jäger, D.; Schleicher, A.; and Westfechtel, B. 1999. Using UML for software process modeling. *ACM SIGSOFT Software Engineering Notes*, 24(6): 91–108.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9): 389–399.
- John, M. M.; Holmström Olsson, H.; and Bosch, J. 2021. Architecting AI deployment: A systematic review of state-of-the-art and state-of-practice literature. In *Software Business: 11th International Conference, ICSOB 2020, Karlskrona, Sweden, November 16–18, 2020, Proceedings 11*, 14–29. Springer.
- Kaminski, M. E. 2021. The right to explanation, explained. In *Research handbook on information law and governance*, 278–299. Edward Elgar Publishing.
- Kaye, K.; Dixon, P.; and Gellman, R. 2023. Risky Analysis: Assessing and Improving AI Governance Tools – An International Review of AI Governance Tools and Suggestions for Pathways Forward. Technical report, World Privacy Forum. EBook/Digital: ISBN 978-0-9914500-2-2.
- Lai, V.; Chen, C.; Smith-Renner, A.; Liao, Q. V.; and Tan, C. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1369–1385.
- Le, T.; Miller, T.; Singh, R.; and Sonenberg, L. 2023. Explaining model confidence using counterfactuals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11856–11864.
- Logg, J. M.; Minson, J. A.; and Moore, D. A. 2019. AI-algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151: 90–103.
- Madden, D.; Cadet-James, Y.; Atkinson, I.; and Watkin Lui, F. 2014. Probes and prototypes: a participatory action research approach to codesign. *CoDesign*, 10(1): 31–45.
- McGregor, S. 2021. Preventing repeated real world AI failures by cataloging incidents: The AI incident database. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 15458–15463.
- McKernon, E.; Glasser, G.; Cheng, D.; and Hadfield, G. 2024. AI Model Registries: A Foundational Tool for AI Governance. *arXiv preprint arXiv:2410.09645*.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.
- Meng, X.-L. 2025. Data Science and AI: Everything Everywhere All at Once. *Harvard Data Science Review*, 7(1).

- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, 220–229.
- Mökander, J.; and Floridi, L. 2023. Operationalising AI governance through ethics-based auditing: an industry case study. *AI and Ethics*, 3(2): 451–468.
- Mokoena, N.; and Obagbuwa, I. C. 2025. An analysis of artificial intelligence automation in digital music streaming platforms for improving consumer subscription responses: a review. *Frontiers in Artificial Intelligence*, 7: 1515716.
- National Archives and Records Administration. 2025. Inventory of NARA Artificial Intelligence (AI) Use Cases. Technical report, U.S. National Archives.
- Noy, C. 2008. Sampling knowledge: The hermeneutics of snowball sampling in qualitative research. *International Journal of social research methodology*, 11(4): 327–344.
- Office of the Federal Chief Information Officer. 2025. Federal AI Use Case Inventory. Technical report, U.S. Federal Government.
- Organisation for Economic Co-operation and Development. 2024. OECD AI Policy Observatory. <https://oecd.ai>. Accessed: 2025-05-23.
- Organisation for Economic Co-operation and Development. 2025a. Tools for Trustworthy AI.
- Organisation for Economic Co-operation and Development. 2025b. Towards a Common Reporting Framework for AI Incidents. OECD Artificial Intelligence Papers, No. 18. Accessed May 19, 2025.
- Parker, C.; Scott, S.; and Geddes, A. 2019. Snowball sampling. *SAGE research methods foundations*.
- Poinet, P.; Stefanescu, D.; Tsakiridis, G.; de Boissieu, A.; and Papadonikolaki, E. 2020. Supporting collaborative design and project management for AEC using Speckle’s interactive data flow diagram. In *Design Computation Input/Output*. Design Computation, London, United Kingdom.
- ProjectManager. 2025. Using a Decision Flowchart in Project & Process Management. Accessed May 20, 2025.
- Sauer, J.; Kao, C.-S.; and Wastell, D. 2012. A comparison of adaptive and adaptable automation under different levels of environmental stress. *Ergonomics*, 55(8): 840–853.
- Schneider, J.; Abraham, R.; Meske, C.; and Vom Brocke, J. 2023. Artificial intelligence governance for businesses. *Information Systems Management*, 40(3): 229–249.
- Shao, Y.; Zope, H.; Jiang, Y.; Pei, J.; Nguyen, D.; Brynjolfs-son, E.; and Yang, D. 2025. Future of Work with AI Agents: Auditing Automation and Augmentation Potential across the US Workforce. *arXiv preprint arXiv:2506.06576*.
- Shortliffe, E. H.; and Buchanan, B. G. 1975. A model of inexact reasoning in medicine. *Mathematical biosciences*, 23(3-4): 351–379.
- Skitka, L. J.; Mosier, K. L.; and Burdick, M. 1999. Does automation bias decision-making? *International Journal of Human-Computer Studies*, 51(5): 991–1006.
- Slattery, P.; Saeri, A. K.; Grundy, E. A. C.; Graham, J.; Noetel, M.; Uuk, R.; Dao, J.; Pour, S.; Casper, S.; and Thompson, N. 2024. The AI Risk Repository: A Comprehensive Meta-Review, Database, and Taxonomy of Risks From Artificial Intelligence. *arXiv preprint arXiv:2408.12622*.
- Slattery, P.; Saeri, A. K.; Grundy, E. A. C.; Graham, J.; Noetel, M.; Uuk, R.; Dao, J.; Pour, S.; Casper, S.; and Thompson, N. 2025. The AI Risk Repository: A Comprehensive Meta-Review, Database, and Taxonomy of Risks From Artificial Intelligence. *arXiv:2408.12622*.
- Steyvers, M.; Tejada, H.; Kerrigan, G.; and Smyth, P. 2022. Bayesian modeling of human-AI complementarity. *Proceedings of the National Academy of Sciences*, 119(11): e2111547119.
- United Nations Educational, Scientific and Cultural Organization. 2024. Global Observatory on Artificial Intelligence Ethics and Governance. <https://globalaiethics.org>. Accessed: 2025-05-23.
- Upmann, P. 2024. AI Escalation Protocols: Addressing Risks with Clear and Efficient Procedures. *AI Governance Institute*.
- Wachter, S.; Mittelstadt, B.; and Floridi, L. 2017. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International data privacy law*, 7(2): 76–99.
- World Economic Forum. 2024. AI Governance Lighthouse Case Study: University of Technology Sydney (UTS). <https://www.weforum.org/publications/ai-governance-lighthouse-case-study-university-of-technology-sydney-uts>. Accessed: 2025-05-23.
- Yonyx. 2024. Explaining AI with the Help of Flowchart Diagrams. Accessed May 20, 2025.