Digital Society (2024) 3:7
https://doi.org/10.1007/s44206-024-00088-0

ORIGINAL PAPER

Governing Ethical Gaps in Distributed AI Development

Nandhini Swaminathan[1] · **David Danks**[2,3]

Received: 11 July 2023 / Accepted: 16 January 2024 / Published online: 15 February 2024

Abstract

Good business practice often leads companies to subdivide into separate functional entities for operational efficiency and specialization. However, these kinds of divisions can generate significant ethical and perhaps even regulatory gaps when they occur in AI companies. In particular, one natural division for an AI company is into separate entities responsible for model development, testing, and cybersecurity (to maintain and protect data). In this paper, we argue that this division can lead to some ethical responsibilities always being "someone else's job." For concreteness, we consider the US National Institute of Standards and Technology's AI Risk Management Framework (NIST AI RMF) as a guide to ethical obligations in a corporate context. We show that a common division of labor in AI development and deployment can lead to specific obligations for which no entity is responsible, even though they apply to the effort as a whole. We propose "Join Accountability Agreements", a mechanism to ensure that ethical obligations do not slip through the cracks because of the way an effort is structured. We thus aim to highlight the significance of comprehensive examinations of and adaptable strategies for our ethical obligations when developing AI systems in a distributed manner.

Keywords NIST AI Risk Management Framework (AI RMF) · AI governance · Risk management strategy · Accountability

✉ Nandhini Swaminathan
nswaminathan@ucsd.edu

David Danks
ddanks@ucsd.edu

[1] Computer Science and Engineering Department, University of California, San Diego, CA, USA

[2] The Halıcıoğlu Data Science Institute, University of California, San Diego, CA, USA

[3] Department of Philosophy, University of California, San Diego, CA, USA

1 Introduction

At a high level, AI governance is an interdisciplinary approach to ensure the ethical and responsible development, deployment, and use of artificial intelligence (AI) technologies. This can involve a range of practices, including creating legal frameworks, adhering to ethical guidelines, and implementing risk management practices (Cihon, 2019; Dafoe, 2018; Perry & Uuk, 2019). The typical goal of AI governance is to align AI systems with human well-being, respect for human autonomy, social responsibility, transparency, and other accountability principles while minimizing adverse effects.

In other technical fields (e.g., cybersecurity), risk assessment and management frameworks have provided valuable tools for governance (Ahmed, 2007). Unsurprisingly, we are now seeing the emergence of AI risk assessment and management frameworks to support *AI* governance (Afzal, 2021; Attard-Frost, 2022; Berk, 2021; Mäntymäki, 2022; ÓhÉigeartaigh et al., 2020; Schmitt, 2022; Taeihagh, 2021). These frameworks aim to provide structured methodologies to identify, assess, and mitigate AI risks, while also promoting various positive features such as transparency, accountability, and sustainability (Schwartz et al., 2022). These frameworks thus require consideration of ethical, legal, and social implications. Many of these AI risk assessment frameworks have been developed by governmental entities, including Australia's AI Assurance framework, the European Commission's Assessment List for Trustworthy Artificial Intelligence (ALTAI), and the Algorithm Impact Assessment tool (AIA) by the Government of Canada (European Commission, 2020; Government of Canada, 2023; UK Information Commissioner, 2022; Chik, 2013; McKelvey and MacDonald, 2019; World Economic Forum, 2022). One particularly prominent approach is the AI Risk Management Framework (RMF) developed by the National Institute of Standards and Technology in the US (AI RMF; Tabassi, 2023a). The AI RMF aims to comprehensively articulate both general principles and goals, and also specific methods and processes to achieve those goals, for socio-technical AI systems. The AI RMF thus aspires to provide a comprehensive approach to manage AI risks, encompassing legal compliance, risk management, and ethical considerations (Schuett, 2022).

The AI RMF provides a structure for risk assessment and management across an entire AI effort, from design through development into deployment and use. However, many AI systems are built in a relatively distributed fashion, with distinct entities—perhaps in the same company, perhaps in different ones—contributing different aspects. For example, one group or company might collect the data, while another does the analysis and model building, while a third sells and deploys the system worldwide. It is thus important that the different parts of risk assessment and management can be assigned to one or more of these entities to ensure that the effort as a whole is ethical. However, we argue that this cannot always be done. We focus below on one particular kind of organizational structure (Sect. 3), but the lessons clearly would apply in other cases. We first begin, though, with a discussion of the AI RMF (Sect. 2) before concluding with a proposal to ensure that ethical responsibilities do not slip through the organizational cracks (Sect. 4).

2 A Primer on the NIST AI Risk Management Framework

The NIST AI Risk Management Framework intends to provide a comprehensive and systematic approach for organizations to navigate the complexities of AI risk management. At its core, the AI RMF comprises four essential functions (which we summarize in Fig. 1): GOVERN, MAP, MEASURE, and MANAGE, each playing a pivotal role in ensuring responsible and effective AI design, development and deployment (Tabassi, 2023b). Each of the high-level functions has a number of (sub-)tasks (omitted from Fig. 1) that operationalize the overall functional goals.

The GOVERN function includes tasks centered on establishing policies, procedures, and practices that should align with the organization's guiding principles and strategic objectives. This foundation should enable the organization to take a proactive approach to AI risk management. Building on this foundation, the MAP function involves carefully identifying and analyzing AI-related risks and their potential ramifications so the organization can hopefully better understand the broader implications of its AI technologies on users and society. The MEASURE function is designed to make the risk management process more rigorous by ensuring that appropriate quantitative and qualitative tools are used to obtain objective, transparent insights for decision-making. Finally, tasks in the MANAGE function help to synthesize insights from the other functions to prioritize and address AI risks, including targeted responses and resource allocations.

Ideally, the tasks in the four NIST AI RMF functions produce a risk management approach that aligns technical design with organizational values in beneficial ways. At the time of this writing, there have not been any systematic studies of the efficacy of the AI RMF in enhancing AI system trustworthiness, optimizing performance, or contributing positively to societal well-being.

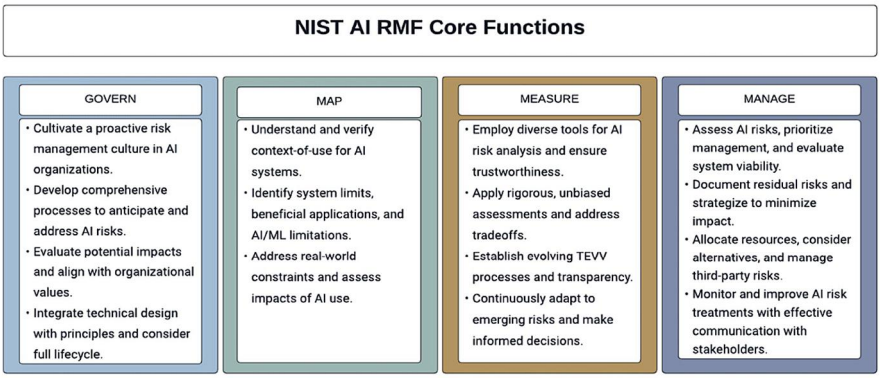

Fig. 1 High-level NIST Risk Management Framework overview: core functions

3 Division of an AI Effort and Its Ethical Implications

In practice, AI efforts often occur in a distributed manner, as different entities have different expertise and capabilities. Functional modules, or distinct units within an AI project each focusing on a specific aspect of development, allow for this distribution of tasks. Although some companies or organizations can perform every step in the lifecycle of an AI effort, they do this through different subgroups (e.g., different divisions within the company) that typically have different roles.[1] By allocating distinct functions to distinct entities, we can reap the benefits of specialization and experience on specific tasks. Figure 2 provides our preferred way of decomposing the lifecycle of an AI system into modular elements, though we note that other divisions could be used instead. The key is that modern AI efforts are almost always pursued in a more modular way, as this approach can promote efficiency and effectiveness within each functional domain through streamlined workflows, optimized resource allocation, and improved performance. Of course, the process is rarely as clean-cut as suggested by Fig. 2: boundaries between functional modules can be blurry; later modules might need to revisit earlier decisions; and so on. Nonetheless, we contend that Fig. 2 is a useful approximation to the practice of modern AI system creation.

One very natural division is into three distinct units: a model development group, a testing group, and a cybersecurity team that ensures the security and integrity of the data and models.[2] These could be three different divisions within the same company, or three different legal entities (e.g., the model development group might be a company that specializes in AI as a Service), or some combination. The model development team constructs and tailors the AI model according to the agreed-upon specifications. The testing unit assesses the AI system's performance and functionality. And the cybersecurity unit safeguards the AI system against threats and attacks. Collectively, these three units play a pivotal role in the overall engineering process, as detailed in Table 1.

However, this separation of responsibilities brings forth its own set of ethical challenges. The AI industry, especially in its early days and even today to some

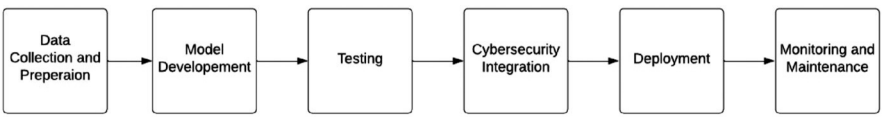

Fig. 2 AI engineering process

[1]Academic efforts are perhaps the exception that proves the rule, as a single small academic research group may have to do everything itself.

[2]We omit the critically important unit for Data Collection because that step largely falls outside of the scope of the NIST AI RMF, which is the framework that we consider for specificity. We emphasize, though, that the inclusion of additional units for Data Collection (or Monitoring, or other functional tasks) would only worsen matters in terms of ethical gaps in risk assessment frameworks. That is, our discussion here is arguably the easiest case for risk assessment (and it is still problematic).

Table 1 Overlap of labor during the engineering process

Process step	Model development team	Cybersecurity team	Testing team
Data collection and preparation	Define data needs, clean the data, etc.	Ensure data sources are secure and data privacy is maintained	–
Model development	Choose algorithms, train model, validate model, etc.	–	Test on evaluation metrics and provide feedback
Testing	Modify the model based on feedback from the testing team	–	Performance testing, robustness evaluation
Cybersecurity integration	–	Vulnerability assessment, integrate security protocols	Ensure security features work as designed
Deployment	Deploy to production, integrate with the application, etc.	Ensure secure deployment environment	Post-deployment checks
Monitoring and maintenance	Performance monitoring, model updates, etc.	Ongoing security monitoring, threat detection	Continuous performance checks, re-testing after updates

extent, has seen cases where important aspects of AI development are contracted out to different entities or segregated within the same organization without comprehensive contractual clarification on responsibility and liability. This is not unique to AI; many nascent industries often face similar challenges. This division of cognitive labor can thus raise significant questions regarding accountability in cases where negative consequences arise from the technology. Determination of (ethical) responsibility can be challenging, as each may argue that its role was limited to its specific function. For instance, suppose the AI system built by these entities exhibits biased or discriminatory behavior. The testing team may argue that its role did not involve developing or implementing the technology, and so it should not be held liable for biases. On the other hand, the model development unit could claim that its responsibility was solely to construct a functional model, shifting the blame elsewhere.

4 Accountability Gaps

4.1 Responsibility Through the Cracks

These efforts, though practical and necessary, create a cascade of challenges that are unaddressed by the NIST AI RMF. These include unclear accountability for AI risks, difficulties in fostering a unified culture of risk awareness, and complexities in standardizing engagement with third parties. Additionally, the compartmentalization may obstruct a holistic view of AI system risks, lead to inconsistent tracking and prioritization of risks, and complicate the documentation and monitoring of risk mitigation strategies. As such, it becomes vital to have strong communication and oversight to counter these challenges. In some cases, there will be legal contracts or internal organizational rules and processes that address

issues such as legal liability. However, those documents (including terms & conditions) essentially never consider ethical risks or liability.

The functions of the NIST AI RMF do not map onto the functional modules through which AI systems are created, and so some of the RMF functions can slip through the cracks with no one assuming ownership. If the use of the NIST AI RMF were legally required, then a regulator (or other governmental entity) could perhaps ensure that nothing slips through the cracks. More generally, if there is a single locus of oversight or coordination for use of a risk assessment framework in AI creation, then one could potentially ensure that no ethical responsibility slips through the cracks. We are not, however, in such a world, and so commitment to using, say, the AI RMF does not thereby ensure that all relevant ethical risks will be considered.

4.2 Joint Accountability Agreements

We propose the establishment of Joint Accountability Agreements (JAA) as a potential solution to address the accountability and responsibility challenges emerging from the scenario of separate entities within an AI effort. In particular, we propose that JAAs can help to ensure accountability when work occurs across multiple entities, taking into account the unique characteristics of AI systems and their societal impacts. Similar ideas for AI accountability have been independently developed (e.g., Berscheid & Roewer-Despres, 2019), though they are not grounded in risk assessment frameworks such as the AI RMF or other mechanisms.

Building upon the existing GOVERN function of the NIST AI RMF, which focuses on key aspects such as legal compliance, risk management, and ethical considerations, we propose that distinct functional modules should establish JAAs between them. These agreements would clearly outline the ethical roles, responsibilities, and accountabilities, including the following components:

1. Roles, responsibilities, and resource allocation
2. Legal and regulatory compliance
3. Agreed-upon framework for risk management
4. Safety, ethical standards, and performance metrics
5. Liability and accountability protocols
6. Monitoring and auditing procedures
7. Termination and transition.

In many ways, a JAA parallels existing legal agreements when work occurs in a distributed fashion (e.g., contracts, spec sheets, etc.), but we propose that these should focus on accountability and responsibility obligations rather than technical or financial ones. In particular, these seven components can be used to ensure that relevant tasks of the GOVERN function are satisfied, even though no single unit has primary responsibility for them (i.e., they would otherwise slip through the cracks):

- GOVERN 1.1 (Legal and Regulatory Requirements): Establish JAA alongside understanding, managing, and documenting legal and regulatory requirements.

- GOVERN 4.1 (Fostering a Safety-first Mindset): Adhere to the agreed-upon standards to mitigate negative impacts.
- GOVERN 6.1 (Third-Party Entity Risks): Define the responsibilities and obligations of third-party entities through JAA to ensure coordination, communication, and accountability.

More generally, in the scenario being discussed, the three entities—the model development team, the testing team, and the cybersecurity team—could decide to divide the NIST AI RMF functions among themselves using JAAs, which would lead to a more comprehensive approach to AI risk governance. This arrangement would allow each entity to focus on its specific areas of expertise, fostering better coordination and a more cohesive risk management strategy, while still ensuring that the overall project is ethical. For instance, the model development entity could oversee most of the GOVERN, MAP, and MEASURE functions. The testing entity could concentrate on the MEASURE and MANAGE functions. The cybersecurity entity could handle portions of the GOVERN and MANAGE functions. This division would ensure that each function gets the right expertise and promotes effective collaboration. It also clarifies roles, fostering accountability within the NIST AI RMF framework. Yet, with this clarity in roles comes the imperative of upholding their respective responsibilities. Any breach of the terms outlined in the JAA could invoke consequences pre-determined by the involved entities, varying from restorative measures for minor infractions to termination of the partnership or even initiation of legal proceedings for major breaches. This acts as a safeguard, ensuring accountability within the NIST AI RMF framework.

At the current time, there are no legal mandates to use particular risk assessment frameworks, though such requirements may be forthcoming. Regardless of the legal requirements, however, JAAs provide a way for modular efforts to ensure that they are satisfying their local ethical obligations within an overall plan that covers all relevant ethical requirements. We thus contend that there is ethical value to JAAs, even if they are not legally required (at the moment). We also emphasize that we have deliberately not specified the exact structure of JAAs (beyond the seven components) as we believe that these will often depend on the particular entities and goals for an AI effort.

4.3 Recommended Solutions for Additional Vulnerabilities

While the establishment of Joint Accountability Agreements (JAAs) could represent a significant step towards addressing accountability and responsibility challenges, there are additional concerns that arise for risk assessment frameworks applied to AI. One issue is the expertise and adherence to the framework required for the involved entities. For example, the testing and cybersecurity modules may not primarily specialize in AI, and so could possess a limited understanding of the NIST AI RMF principles and practices, resulting in suboptimal implementation. This discrepancy could lead to inconsistent application of risk management strategies and potential oversight of critical vulnerabilities. One response would look to voluntary training and certification, as they could help to ensure that all teams

involved have appropriate knowledge bases. A different (not mutually exclusive) response would be to better characterize the knowledge, skills, and processes required for the different AI RMF (sub-)tasks so that the functional modules can ensure that they have the needed expertise.

Additionally, the lack of common standards or protocols for AI technologies presents a notable challenge, particularly in industries that rely on multi-vendor and multiplatform AI solutions. Without standardization, interoperability (including ethical interoperability; Danks & Trusilo, 2022) between AI systems may be compromised, resulting in inefficiencies and diminished effectiveness. Furthermore, non-standardized AI technologies can introduce security vulnerabilities and unintended consequences that can impact stakeholders and overall system performance. For instance, in the context of predictive maintenance in industrial settings, the absence of common standards may impede the seamless integration of predictive models from different vendors, thereby reducing maintenance effectiveness and increasing the risk of equipment failure. Encouraging the adoption of common standards within the NIST AI RMF can facilitate interoperability, enhance system security, and promote a harmonized approach to AI risk management across industries.

5 Conclusion

In our paper, we examine the challenges and potential issues arising from dividing an AI effort into separate entities responsible for model development, cybersecurity, and testing. We found that such a division creates challenges in coordination, communication, responsibility allocation, and the possibility of overlooking critical vulnerabilities due to fragmented oversight.

We have focused on the NIST AI RMF to provide specification about these concerns and proposed incorporating Joint Accountability Agreements (JAAs) into the GOVERN function. JAAs aim to improve AI risk governance and accountability among the involved entities by ensuring that all parties share responsibility for the AI system's performance, ethical alignment, and risk mitigation strategies. The integration of JAAs can foster collaboration by establishing a shared understanding of the entities' roles, responsibilities, and expectations.

At the same time, our focus on the AI RMF inevitably means that some issues have received less attention. For example, the AI RMF is largely focused on analysis of data rather than the data itself. However, there is increasing awareness of the importance of data for the performance (including ethical implications) of AI systems, particularly large language models and other data-intensive models. Ethical responsibilities could fail to be met because of the division of labor between data collectors and data modelers, and so JAAs could also prove useful in this case, even if they are not required to meet the demands of the AI RMF.

Furthermore, JAAs are arguably important and valuable across a range of other AI risk frameworks like the European Union's ALTAI, Australia's AI Assurance

Framework, and Canada's Algorithmic Impact Assessment (AIA). Each of these frameworks emphasizes somewhat different aspects: ALTAI focuses on trustworthiness, Australia's framework emphasizes ethical principles and self-assessment, and Canada's AIA prioritizes public engagement and transparency. However, none of them address the nuances that emerge from multi-entity collaborations; in particular, they do not have mechanisms to ensure that ethical responsibilities do not slip through the cracks.

Our analysis highlights the importance of establishing clear accountability and responsibility among separate entities involved in AI development. By proposing the integration of JAAs alongside the NIST AI RMF, we contribute to the ongoing conversation on AI risk management and governance, emphasizing the need for comprehensive and adaptable strategies that align with ethical standards and societal values. However, we acknowledge that implementing JAAs requires a high degree of trust and transparency between entities, which may not always be feasible in competitive environments or when dealing with sensitive information. Furthermore, the effectiveness of JAAs depends on the willingness and capacity of the entities to collaborate and hold each other accountable. Future research should explore alternative mechanisms for fostering accountability and responsibility in such situations.

As the AI landscape continues to evolve rapidly, it is crucial to develop comprehensive and adaptable AI governance frameworks that address such unique challenges and ensure that AI systems contribute positively to societal well-being and adhere to ethical principles. At the same time, we must work to ensure that the distributed nature of AI system creation does not create cracks through which ethical responsibility can slip.

Author Contributions N.S. devised the project, the main conceptual ideas, and the proof outline. The first draft of the manuscript was written by N.S. D.D. supervised the project. Both N.S. and D.D. contributed to the final version of the manuscript. All authors read and approved the final manuscript.

Funding Not applicable.

Data Availability Not applicable.

Declarations

Informed Consent Not applicable.

Competing Interests Not applicable.

References

Afzal, F., Yunfei, S., Nazir, M., & Bhatti, S. M. (2021). A review of artificial intelligence based risk assessment methods for capturing complexity-risk interdependencies: Cost overrun in construction projects. *International Journal of Managing Projects in Business, 14*(2), 300–328. https://doi.org/10.1108/IJMPB-02-2019-0047

Ahmed, A., Kayis, B., & Amornsawadwatana, S. (2007). A review of techniques for risk management in projects. *Benchmarking: An International Journal, 14*(1), 22–36. https://doi.org/10.1108/14635770710730919

Attard-Frost, B., De Los Rıos, A., & Walters, D. R. (2022). The ethics of AI business practices: A review of 47 AI ethics guidelines. *AI and Ethics,* 1–18. https://doi.org/10.2139/ssrn.4034804

Aziz, S., & Dowling, M. (2019). Machine learning and AI for risk management. In T. Lynn, J. G. Mooney, P. Rosati, & M. Cummins (Eds.), *Disrupting finance: FinTech and strategy in the 21st century* (pp. 33–50). Springer International Publishing. https://doi.org/10.2139/ssrn.3201337

Berk, R. A. (2021). Artificial intelligence, predictive policing, and risk assessment for law enforcement. *Annual Review of Criminology, 4,* 209–237. https://doi.org/10.1146/annurev-criminol-051520-012342

Berscheid, J., & Roewer-Despres, F. (2019). Beyond transparency: A proposed framework for accountability in decision-making AI systems. *AI Matters, 5*(2), 13–22. https://doi.org/10.1145/3340470.3340476

Chik, W. B. (2013). The Singapore Personal Data Protection Act and an assessment of future trends in data privacy reform. *Computer Law & Security Review, 29*(5), 554–575. https://doi.org/10.1016/j.clsr.2013.07.010

Cihon, P. (2019). Standards for AI governance: International standards to enable global coordination in AI research & development. Future of Humanity Institute. University of Oxford.

Dafoe, A. (2018). *AI Governance: A research agenda* (Vol. 1442, p. 1443). Governance of AI Program, Future of Humanity Institute, University of Oxford.

Danks, D., & Trusilo, D. (2022). The challenge of ethical interoperability. *Digital Society, 1,* 11.

European Commission. (2020). *The assessment list for trustworthy artificial intelligence*. European Commission High-Level Expert Group on Artificial Intelligence. https://altai.insight-centre.org/

Garvey, C. (2018). AI risk mitigation through democratic governance: Introducing the 7-dimensional AI risk horizon. In *Proceedings of the 2018 AAAI/ACM conference on AI, ethics, and society* (pp. 366–367). https://doi.org/10.1145/3278721.3278801

Government of Canada. (2023). *Algorithmic impact assessment tool*. Government of Canada. https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html

Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics, 2*(4), 603–609. https://doi.org/10.1007/s43681-022-00143-x

McGilvray, D. (2021). *Executing data quality projects: Ten steps to quality data and trusted information* (TM). Academic Press.

McKelvey, F., & MacDonald, M. (2019). Artificial intelligence policy innovations at the Canadian federal government. *Canadian Journal of Communication, 44*(2), PP–43. https://doi.org/10.22230/cjc.2019v44n2a3509

Model AI governance framework. (n.d.). Personal Data Protection Commission Singapore. https://www.pdpc.gov.sg/help-and-resources/2020/01/model-aigovernance-framework

Nsw artificial intelligence assurance framework. (n.d.). Australia NSW Government. https://www.digital.nsw.gov.au/policy/artificial-intelligence/nsw-artificialintelligence-assurance-framework

ÓhÉigeartaigh, S. S., Whittlestone, J., Liu, Y., Zeng, Y., & Liu, Z. (2020). Overcoming´ barriers to cross-cultural cooperation in AI ethics and governance. *Philosophy and Technology, 33,* 571–593. https://doi.org/10.1007/s13347-020-00402-x

Perry, B., & Uuk, R. (2019). AI governance and the policymaking process: Key considerations for reducing AI risk. *Big Data and Cognitive Computing, 3*(2), 26. https://doi.org/10.3390/bdcc3020026

Rfd bus012a artificial intelligence assessment tool. (n.d.). US Pennsylvania Office of Administration. https://www.oa.pa.gov/Policies/Documents/rfd-bus012a.xlsx

Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. (2021). "Everyone wants to do the model work, not the data work": Data cascades in high-stakes AI. In Proceedings of

the 2021 CHI conference on human factors in computing systems (pp. 1–15). ACM. Held virtually; originally Yokohama, Japan, May 8–13. https://doi.org/10.1145/3411764.3445518

Schmitt, L. (2022). Mapping global AI governance: A nascent regime in a fragmented landscape. *AI and Ethics, 2*(2), 303–314. https://doi.org/10.1007/s43681021-00083-y

Schuett, J., & Anderljung, M. (2022). Comments on the initial draft of the NIST AI risk management framework.

Schwartz, R., Vassilev, A., Greene, K., Perine, L., Burt, A., & Hall, P. (2022). Towards a standard for identifying and managing bias in artificial intelligence. *NIST Special Publication, 1270,* 1–77.

Tabassi, E. (2023a). Artificial intelligence risk management framework (AI rmf 1.0).

Tabassi, E. (2023b). *Artificial intelligence risk management framework playbook*. https://airc.nist.gov/AI_RMF_Knowledge_Base/Playbook

Taeihagh, A. (2021). Governance of artificial intelligence. *Policy and Society, 40*(2), 137–157. https://doi.org/10.1080/14494035.2021.1928377

UK Information Commissioner. (2022). *AI and data protection risk toolkit*. UK Information Commissioner's Office. https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/

World Economic Forum. (2022). *Artificial intelligence for children toolkit*. World Economic Forum. https://www3.weforum.org/docs/WEF-Artificial-Intelligencefor-Children-2022.pdf