Conference Proceedings

DOUGLAS YEUNG, TINA HUANG, BENJAMIN BOUDREAUX, PRATEEK PURI, JONATHAN W. WELBURN, ANITA CHANDRA, MIRIAM VOGEL

Steps Toward AI Governance

Insights and Recommendations from the 2024 EqualAI Summit

For more information on this publication, visit **www.rand.org/t/CFA3799-1**.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.
© 2025 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

About These Conference Proceedings

EqualAI's 2024 Responsible AI Summit, cosponsored by RAND, was convened on July 31, 2024, in Washington, D.C., to facilitate dialogue among corporate stakeholders from multiple industries, functions, and roles about artificial intelligence (AI) development, acquisition, and integration. The purpose of the summit was to identify and align on common practices, discuss challenges, and share lessons learned in establishing and evaluating metrics in AI governance.

In these conference proceedings, we describe key insights derived from summit discussions about best practices, metrics, and tools for evaluating the standards and performance of AI systems. These insights are intended to help organizations foster a cohesive approach to AI governance.

Community Health and Environmental Policy Program

RAND Social and Economic Well-Being is a division of RAND that seeks to actively improve the health and social and economic well-being of populations and communities throughout the world. This research was conducted in the Community Health and Environmental Policy Program within RAND Social and Economic Well-Being. The program focuses on such topics as infrastructure, science and technology, industrial policy, community design, community health promotion, migration and population dynamics, transportation, energy, and climate and the environment, as well as other policy concerns that are influenced by the natural and built environment, technology, and community organizations and institutions that affect well-being. For more information, email chep@rand.org.

RAND Global and Emerging Risks

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. For more information, visit www.rand.org/global-and-emerging-risks or contact the division director (contact information is provided on the webpage).

EqualAI

EqualAI is a nonprofit organization that helps companies, policymakers, and leading institutions define and implement AI governance frameworks that foster innovation and enable broader AI adoption.

Funding

Funding for this work was provided by gifts from RAND supporters and EqualAI.

Acknowledgments

We are grateful for the individuals who, as members of the EqualAI Badge Program, shared their insights before and during the summit. We also thank those who supported this effort, including Natalie Heinitz, Teddy Tawil, and Jim Wiley at EqualAI and Heather Gillispie at RAND. We also thank our peer reviewers, Neeti Pokhriyal and Osonde Osoba, for their valuable feedback.

Contents

Steps Toward AI Governance: Insights and Recommendations from the 2024 EqualAI Summit

Over the past few years, large companies have been early adopters of artificial intelligence (AI),[1] and the technology has increasingly permeated various sectors, driving transformative changes across industries and society at large. As the deployment of AI technologies accelerates, the need for effective AI development and use becomes ever more pressing. In 2023, EqualAI's second AI summit convened government and industry leaders in a significant effort to discuss these needs and advance alignment on metrics for an AI governance framework applicable to organizations of any size, industry, and maturity level.[2]

In 2024, EqualAI's third AI summit again convened a diverse array of senior executives—including many from such Fortune 500 companies as Amazon Web Services, Google DeepMind, PepsiCo, Salesforce, Starbucks, and Verizon, as well as other companies, such as LivePerson, M&T Bank, and SAS Institute—who are members of the EqualAI Badge Program.[3] Although many of these organizations are becoming more aware of the challenges associated with AI development, acquisition, and integration, many continue to operate in silos, developing independent AI governance processes and metrics. This fragmented approach underscores the urgent need for information-sharing so organizations can better navigate the challenges and opportunities that arise when formulating a cohesive strategy for effective AI deployment.

Building on discussions from the second summit, the third summit aimed to identify the

- landscape of metrics, tools, evaluation methods, or processes that companies use for AI development and deployment
- technical and organizational challenges that industry leaders face when implementing AI principles, best practices, and an overall governance framework
- solutions to address AI governance challenges that participants recommended for government and industry.

EqualAI and RAND jointly developed these summit objectives and the discussion agenda and then analyzed the results (see Appendix A for more detail). This document highlights key

[1] Brian Eastwood, "The Who, What, and Where of AI Adoption in America," MIT Sloan School of Management, February 7, 2024.

[2] Laxmi Anish, Kathy Baxter, Dawn Bloxwich, Catherine Goetz, Tina Huang, Gabrielle Kohlmeier, David Lincoln, David Marcos, Amanda Muller, Mike Tang, Miriam Vogel, et al., *An Insider's Guide to Designing and Operationalizing a Responsible AI Governance Framework*, EqualAI, September 2023

[3] Throughout these conference proceedings, we refer to these members collectively as the *EqualAI Badge Community* (EqualAI, "EqualAI Badge Program," webpage, undated).

takeaways from discussions at the 2024 EqualAI summit and is intended to serve as a resource for organizations navigating the implementation of effective AI governance.

Existing Approaches to Operationalizing AI Governance Best Practices

A primary objective of the summit was to map the existing landscape of tools, metrics, and methodologies employed for AI governance. Organizations that have adopted AI in various contexts have adapted resources developed by the government, civil society, and industry, tailoring their processes to implement best practices.

When discussing the ways that they mitigate the potential risks associated with AI systems, participants highlighted their reliance on a broad set of AI principles and practices that have emerged in recent years. Many of these approaches are applicable across different sectors and organizations of varying sizes. For instance, the framework developed by EqualAI comprises six key pillars that seek to provide a comprehensive approach to AI governance. Additionally, government agencies have established principles for AI use, such as the Department of Defense's Ethical Principles for AI.[4]

While useful, these sets of broad principles might not help organizations understand how to implement specific governance controls or processes. Therefore, both government and industry have developed more-practical resources and tools to assist organizations in operationalizing AI governance principles. One notable example is the AI Risk Management Framework (RMF) developed by the National Institute of Standards and Technology (NIST). The RMF is a voluntary framework created through a collaborative public-private partnership designed to integrate trustworthiness characteristics into the design, development, and deployment of AI systems. Although the RMF does not prescribe detailed normative standards, it provides key procedural dimensions for robust, effective AI implementations.[5] Furthermore, the private sector has introduced model or system cards to encourage transparent documentation of how AI models are designed, trained, and evaluated.

Several EqualAI Badge Community members shared that they had crafted their own approaches to assessing AI risks and implementing governance. These approaches include (1) enterprise-specific evaluations of algorithmic fairness that use high-level principles to conduct more-nuanced assessments tailored to particular applications and (2) individualized methods of risk tiering and prioritization for the testing and monitoring of AI tools and products.

---

[4] C. Todd Lopez, "DOD Adopts 5 Principles of Artificial Intelligence Ethics," U.S. Department of Defense, February 25, 2020.

[5] National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, U.S. Department of Commerce, January 2023.

Technical Limitations and Constraints on Effective AI Governance

External Models Can Be Difficult to Evaluate

The EqualAI Badge Community agreed that one of the most common technical challenges companies face when evaluating and testing AI models is not knowing how rigorously external models have been evaluated. Specifically, externally acquired models might have undergone different testing and evaluation from those developed in-house. This knowledge gap about external models stems from a lack of transparency regarding their development process, including their initial design and the data they were trained on. Furthermore, participants noted that external models would not have been tested against their organization's use cases nor adapted to their needs and preferences. In contrast, customized testing and adjustments would have been possible for internally developed models or AI models procured and deployed after modifications, such as fine-tuning.

Prioritizing Evaluation of High-Risk Use Cases

Differing use cases and risk levels present a significant technical challenge for companies in prioritizing model testing and evaluations. Organizations must carefully identify which use cases pose the highest risks and allocate their limited resources for testing and evaluation accordingly. While the definition of high-risk use cases varies, the EqualAI Badge Community broadly agreed that the highest risks include scenarios that could lead to bodily harm, financial loss, release of sensitive nonpublic data, human resources (HR) violations, illegal outcomes, or denied opportunities. Defensive applications of AI, such as those used for cybersecurity, are viewed as a distinct category. Additionally, participants believed that generative AI could carry particularly high risks, necessitating more-rigorous assessments of such use cases.

However, even if risk has been properly profiled, it might still be difficult to develop standardized evaluation metrics that capture the vulnerabilities of a generative AI system, given the diversity and stochasticity of its outputs. Participants pointed toward more-flexible evaluation frameworks, such as red teaming or active user feedback, as mechanisms to identify failure modes from generative systems that traditional metrics might not capture.

Organizational Limitations and Constraints on Effective AI Governance

Misaligned Incentives Might Slow or Impede AI Governance

Participants highlighted how misaligned organizational goals can create disincentives for investing the significant resources required to implement appropriate AI processes. For instance, sales and engineering teams might feel pressure to deliver products quickly, making it challenging to integrate AI best practices, such as documentation for transparency, into an already high-pressure environment. Gaining buy-in for AI governance initiatives is particularly

difficult when these practices are not legally mandated. Additionally, there are few incentives for employees to voice concerns, and whistleblower protections are lacking. In such an environment, risks can be overlooked. A changing regulatory landscape, however, could significantly shift this risk environment. For example, new state-level or international regulations could alter incentive structures for AI governance. As new regulations emerge, research will be needed to understand their potential impact.

Company Culture Sets the Tone

Company culture is a key component of AI governance. The EqualAI Badge Community emphasized that cultivating a supportive culture is essential for upholding AI principles, requiring time and commitment to establish effective norms. Shifting the perception of AI governance from that of an unnecessary burden to an essential asset is crucial for teams to achieve their goals efficiently.

Teams must recognize the value of AI governance, including strengthening business partnerships, preventing costly errors, fostering thought leadership, and attracting top talent with relevant skills. Conversely, a lack of effective AI governance could result in legal liabilities, reputational penalties, and/or financial penalties.

The EqualAI Badge Community agreed that an AI governance team needs authority to have impact. As they put it, one factor that contributes to the perception that AI governance teams are "not critical" is when they lack "no-go" authority or a "kill switch" to halt product releases or ongoing processes. Typically, concerns must be validated by legal and compliance teams who do possess such authority before action is taken to slow or stop operations. In general, the authority to restrict a product release might need a clear, impending harm to justify the cost to agility, and the absence of AI-specific laws and regulations in the United States and elsewhere might make it harder to demonstrate these harms.

Leadership Buy-In Is Essential for Sustainable Governance

Leadership support is crucial for fostering an organizational culture that embraces AI governance. While initiatives can start from the bottom up, true sustainability hinges on buy-in from top leadership. This support should extend beyond a single leader to encompass the entire board and C-suite executives. Moreover, given the potential for leadership turnover, securing buy-in from multiple stakeholders is essential for long-term sustainability.

Employee Buy-In Is Crucial

According to participants, a significant challenge in adopting AI governance can stem from employee fears of job displacement, which could foster a culture of hesitance toward AI adoption. This resistance could be exacerbated by a lack of AI literacy throughout the organization. Training employees on AI is challenging because training is likely most effective when provided to those who specifically need such training to do their jobs, such as those who

use AI tools or who make strategic decisions about investing in AI technology. Measuring the success of any AI training presents a related challenge and starts with understanding what level of AI literacy is needed for a given role.

Increasing employee comfort with AI is also dependent on organizational goals and incentives. If there are few opportunities for employees to learn about AI governance and best practices, this could contribute to a lack of consensus about how AI systems should be implemented or governed.

Vendor Relations Should Be Strengthened

Productive vendor relationships can be hindered by an overall lack of technical knowledge about a given AI model (e.g., model architecture, testing and validation results). Procurement teams might not know which technical questions to ask about AI governance, while vendors might be ill-equipped to provide the necessary insights on such crucial aspects as metrics and model transparency. The risk that a generative AI system poses, for example, might depend heavily on the use case and environment in which it is deployed, further complicating matters. Therefore, vendors could be caught in situations in which they must develop a suite of metrics that are applicable to consumers across business domains while being specific and actionable enough for the customized use cases and risk profiles of an individual customer. This gap in foundational understanding and consistent inquiry leads to miscommunication and missed opportunities for ensuring effective AI development and testing.

Recommendations

The EqualAI Badge Community proposed several solutions to address the challenges of implementing AI governance, encompassing recommendations for both companies and government bodies.

Recommendations for Companies

Catalog AI Use Cases

Companies should maintain a centralized catalog of AI tools and their applications, and the catalog should be regularly updated to track use across the organization. This catalog should document key specifications and facilitate risk assessment, prioritizing use cases based on such factors as potential financial loss, HR violations, illegal outcomes, or bodily harm. By clarifying uses and their risk levels, organizations can promote transparency and informed decisionmaking regarding AI.

Standardize Vendor Questions

Developing a standardized set of relevant questions for vendors would enable companies to align on vendor evaluations that are based on standardized metrics. Vendors would be aware

of—and incentivized to meet—key expectations. This approach can facilitate better integration of aligned practices into vendor model development, promoting AI governance across business relationships.

Create an AI Information Tool

Implementing an internal information tool, such as a chatbot, that would provide a company's employees with clear answers to AI governance questions could serve as a valuable resource. This chatbot could be trained with relevant information from government, industry, and nonprofit sources and could be tailored to the company's needs and use cases.

Foster Multistakeholder Engagement

Companies should prioritize both internal and external engagement regarding their AI practices. Internally, this involves securing C-suite support and fostering a culture of AI governance through education and open communication. Externally, organizations should solicit feedback from a broad set of stakeholders, particularly unanticipated end users and those from historically marginalized communities, to understand and mitigate the impacts of AI deployments that might not have been captured in the development process. Additionally, thoroughly documenting these actions throughout the AI life cycle will enhance transparency and aid in cross-enterprise and regulatory discussions.

Leverage Existing Processes

Organizations can apply such established processes as crisis management to handle AI-related incidents (e.g., data leaks) and technical risk management (e.g., code reviews, stress testing) to support AI governance. This approach encourages efficiency by integrating AI governance practices into existing frameworks rather than creating new ones.

Recommendations for the Federal Government

Promote a Consistent Approach

According to summit participants, the federal government should legislate, foster, or incentivize the establishment of a consistent regulatory framework that would help organizations navigate the complexities and standardize the management of AI governance across various states and international jurisdictions. This might involve clarifying existing laws and providing new regulations and laws that govern the use of AI tools specifically. The EqualAI Badge Community highlighted a desire for governments to balance swift regulatory action with the promotion of innovation as AI technologies continue to evolve.

Conclusion

The 2024 EqualAI summit underscored the critical importance of fostering a collaborative environment for effective AI governance. As the adoption of AI technologies accelerates, organizations must move beyond isolated efforts to implement, evaluate, and manage their AI systems. Going forward, organizations should engage in shared learning to effectively tackle the complex AI-related challenges that they face. Summit discussions highlighted a broad array of approaches, tools, and frameworks that are being employed but also revealed significant gaps and limitations in technical evaluation, organizational culture, and employee engagement.

To appropriately navigate the fast-moving landscape of AI, organizations need to prioritize building a culture that values best practices and is supported by leadership across all levels. This effort entails investing in AI literacy, standardizing vendor expectations, and aligning organizational goals with AI governance principles. By fostering an environment in which knowledge and best practices can be shared, businesses can not only mitigate risks but also unlock the full potential of AI technologies.

Looking ahead, the insights from the third EqualAI summit offer considerations for organizations aiming to implement effective structures for AI governance. This imperative is not merely a regulatory need but a strategic imperative that can enhance innovation, build trust, and foster sustainable outcomes. Key recommendations from the summit can support best practices to incorporate and standardize in companies across the country and globe, such as aligning vendor questions, creating an AI information tool, and maintaining a centralized catalog of AI use cases. Fostering multistakeholder engagement and leveraging existing organizational processes on which to establish AI governance practices are key elements of ensuring long-term success in this critical effort.

Building a set of resources or a community to share best practices and lessons learned is, thus, critical to sustaining AI governance efforts. Several summit participants mentioned that, from their perspectives, one of the most-valuable experiences was the opportunity to meet and share insights with peers working on corporate AI governance. The journey toward widespread effective use of AI is ongoing, and collaboration across disciplines and industry sectors will be essential in shaping a future in which AI serves as a positive force for all of society.

Appendix A. Methods

The 2024 EqualAI summit was designed to better understand the technical, organizational, and cultural challenges that enterprises face when operationalizing AI governance and to identify opportunities and solutions for industry and government to mitigate such challenges. With this goal in mind, RAND and EqualAI collaborated to develop summit objectives and a discussion agenda. The summit was a collaborative endeavor featuring the EqualAI Badge Community, EqualAI staff, partners, and RAND researchers focused on technology policy and ethics. Following the summit, RAND and EqualAI collaborated to summarize the key themes and takeaways.

Developing the Agenda

We developed the summit agenda to collect information related to our key objectives. We developed general discussion questions to first allow participants to respond broadly and more-specific questions to then probe specific aspects of those responses. The questions were partially predetermined; they began as targeted and structured questions but were subsequently adapted to the flow of the conversation. We also developed two hypothetical narrative scenarios to encourage participants to explore how their companies might handle a real-world concern related to their use and governance of AI. The full set of questions and scenarios appears in Appendix B. To refine the agenda and initiate discussion, we invited participants to take a presummit survey to help us understand how they were using best practices and principles of AI governance.

This agenda was intended to be flexible, industry agnostic, and broadly applicable to different participants with various roles and responsibilities in their organizational structures. The different types of discussion (e.g., large- versus small-group) and types of questions (e.g., specific versus narrative) were also intended to elicit a greater variety of participation and information.

At the Summit

During the summit, EqualAI shared the presummit survey question results, such as what AI guidance (e.g., NIST RMF) companies found useful, to provide context for discussions. Then, participants split into four (three in-person, one virtual) breakout sessions to first discuss specific topics and then consider a hypothetical scenario. Afterward, all participants reconvened to share and discuss their findings.

Analyzing Themes and Identifying Takeaways

We took several approaches to analyzing discussion themes. First, summit participants jointly discussed the results of their small-group findings. We consolidated these themes in real time, which allowed for additional participant feedback. After the summit, we used the resulting set of themes as a foundation and then discussed additional areas to include in the final analysis reported here.

Our intent was to capture a broad variety of approaches to AI governance, including existing processes and proposed solutions. We did not explicitly analyze the prevalence of given themes. Rather, we used the large-group discussion to elicit themes that garnered broad agreement. In our analysis, we used those broad themes as an organizing framework to categorize practices and solutions. However, we attempted, where possible, to provide some indication of practices or solutions that garnered more agreement. Also, the themes, practices, and solutions represent only the views or statements of the participants; we did not attempt to verify or otherwise evaluate the prevalence of AI governance practices or the effectiveness of proposed solutions.

Appendix B. Discussion Questions and Narrative Scenarios

This appendix contains the full set of discussion questions and scenarios jointly developed by RAND and EqualAI prior to the 2024 EqualAI summit.

Morning Breakout Session 1: Constraints and Obstacles

- We know that algorithms and automation have played a pivotal role in business processes for many years. Today we are most interested in the responsibility challenges and opportunities posed by the more recent boom in AI technologies.
  - Can you describe the most-impactful ways AI has been integrated into your organization's products or processes within the last year? What are the specific products or processes that have been impacted, and what specific role is AI playing?
- What steps need to happen before an AI system is deployed into production? What key personnel would need to be involved?
- Now that we have a better sense of how people are leveraging AI technologies, we want to shift gears to unpacking AI responsibility frameworks.
- What constraints or obstacles has your company faced in aligning with responsible AI principles and/or best practices? We are going to cover this in three dimensions:
  - What *technical* challenges have you faced?
    - **Relevant Participants**: Technical members
  - What *organizational* challenges have you faced?
    - **Relevant Participants**: Management members
  - What *incentivization* or *legal* challenges have you faced?
    - **Relevant Participants**: Legal members

(Proceed to tackle each one by one, starting with technical and working down toward incentives.)

- In light of the challenges that were discussed, what solutions have you explored (or potentially already implemented) to address them?
- What factors are limiting the proposed solutions from being even more effective?
- What additional support from government, researchers, or other actors would be useful in overcoming these challenges?
- What have been the most-helpful concepts or frameworks you've adopted from published Responsible AI materials, such as the NIST AI RMF?
- What concepts, although helpful on a general level, have been most difficult to operationalize? What has been the root of these difficulties?

Morning Breakout Session 1: Hypothetical Scenario

Users are reporting toxic content from a product your organization owns that is powered by generative models. Who would be the first person in your organization to learn about this issue? For that person, who would come to mind as the next people in the organization to involve?

Afternoon Breakout Session 2: Evaluations

- Building from last year's session, participants identified at least one part of Responsible AI as developing monitoring and evaluation systems to pinpoint behaviors of concern before they reach critical mass. We want to dig deeper into how you and your teams are thinking about evaluation.
- Specifically, how do you evaluate models for safety, equity, and security?
  - *Sub-question*: How does the evaluation process change when you are serving clients that have differing levels of comfort with AI analysis pipelines?
- How does your company prioritize which models should be subject to more-rigorous forms of evaluation and testing?
- What types of models does your company currently prioritize for evaluation, and what use cases do these models serve?
- What evaluation methods or strategies do you apply to these systems?
- Have you noticed breaking points in these evaluations or types of unwanted behavior that are difficult to encapsulate into metrics?
- What metrics do you use, if any, to determine whether your evaluations are achieving their intended goal?
  - **Relevant Participants**: Technical members
- What systems and/or processes are in place to operationalize accountability for Responsible AI outcomes?

Afternoon Breakout Session 2: Hypothetical Scenario

An engineer discovers a potential flaw in an AI-powered product pipeline that they believe may cause unexpected outputs, and they are threatening to go public. What is the chain of communication by which this engineer's insight is communicated throughout the organization? Before going public, the engineer develops a metric that can monitor for the occurrence of this behavior in the future. They estimate it will cost $X million to integrate the measurement of this metric into all relevant systems and say they won't go public if it is implemented. What are the first questions that come to mind or people you would think to involve in this conversation?

Abbreviations

AI	artificial intelligence
HR	human resources
NIST	National Institute of Standards and Technology
RMF	Risk Management Framework

About the Authors

Douglas Yeung is an associate director of the Management, Technology, and Capabilities Program in the RAND Homeland Security Research Division and a senior behavioral and social scientist. His research explores how insights about emerging technologies can be used in policymaking. He holds a Ph.D. in psychology.

Tina Huang was the director of strategic initiatives at EqualAI, where she developed key partnerships and scaled EqualAI's programs, such as the Corporate Badge Program and congressional Responsible AI Policy Workshop Series. Huang contributed to these conference proceedings while she was part of EqualAI. She holds an M.A. in security studies.

Benjamin Boudreaux is a policy researcher at RAND working at the intersection of human security, technology, and ethics. His research focuses on policy related to AI, cyberspace, and social media. He holds a Ph.D. in philosophy.

Prateek Puri is an associate information scientist at RAND. His research is aimed at identifying AI threats, assessing their relative degrees of concern to different stakeholders, and developing AI governance frameworks to address them. He holds a Ph.D. in experimental atomic, molecular, and optical physics.

Jonathan W. Welburn is a senior operations researcher at RAND. His research interests include systemic risks, market failures, cybersecurity and technology threats, and deterrence. He holds a Ph.D. in industrial engineering.

Anita Chandra is vice president and director of RAND Social and Economic Well-Being and a senior policy researcher at RAND. Her research explores systems change, program and policy design and implementation, and measure development in a variety of areas. She holds a DrPH.

Miriam Vogel is the president and CEO of EqualAI, and she serves as chair of the National AI Advisory Committee. She has extensive experience working with C-suite executives, boards of directors, policymakers, lawyers, and other key stakeholders to establish best practices for legal and regulatory compliance. She holds a J.D.

References

Anish, Laxmi, Kathy Baxter, Dawn Bloxwich, Catherine Goetz, Tina Huang, Gabrielle Kohlmeier, David Lincoln, David Marcos, Amanda Muller, Mike Tang, et al., *An Insider's Guide to Designing and Operationalizing a Responsible AI Governance Framework*, EqualAI, September 2023. As of February 3, 2025: https://www.equalai.org/wp-content/uploads/2023/09/EQUAL-AI_Whitepaper.pdf

Eastwood, Brian, "The Who, What, and Where of AI Adoption in America," MIT Sloan School of Management, February 7, 2024.

EqualAI, "EqualAI Badge Program," webpage, undated. As of January 24, 2025: https://www.equalai.org/programs/equalai-badge-program/

Lopez, C. Todd, "DOD Adopts 5 Principles of Artificial Intelligence Ethics," U.S. Department of Defense, February 25, 2020.

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, U.S. Department of Commerce, January 2023.