

# EXECUTIVE SUMMARY

As AI, machine learning algorithms, and algorithmic decision systems (ADS) continue to permeate every aspect of our lives and our society, the question of AI governance becomes exceedingly important. From racially biased healthcare algorithms to AI-enabled targeting decisions and from opaque and biased hiring algorithms to self-driving cars, the potential for AI and ADS to cause harm and infringe on both individual and group rights is significant. This is why increasingly more regulations are being proposed to audit or evaluate the impacts of algorithms that make or contribute to morally and legally consequential decisions. Alongside this increase in regulation, there has been a significant uptick in interest regarding the internal governance of AI. Organizations and institutions both large and small, nonprofit and for-profit, private and public, have begun creating and implementing governance tools and structures to ensure: (a) compliance with upcoming regulation, (b) minimization of reputational and financial risks of bad algorithms, and (c) safety and adherence to ethical standards for the responsible use of AI. This report examines the current state of internal governance structures and tools across organizations, both in the private and public sectors and in large and small organizations. This report provides one of the first robust and broad insights into the state of AI governance in the United States and Europe.

Using a literature review, survey, and interviews we interrogated three related questions. We asked: (1) what governance tools are being used for AI across different sectors and industries, (2) what tools seem to be working, and (3) why are they working? The result is a robust and comprehensive picture of the range of governance tools and practices organizations use and a preliminary assessment of whether those tools and practices work for those organizations.

Our analysis found that significantly less than half of all organizations that use or develop AI have any formal or substantial governance structures for AI. Among those that do have intentional AI governance structures, there is a variety of governance tools being used and a variety of reasons for adopting them. This general heterogeneity in the field is likely a result of no clearly established best practices or standards for AI governance.

“  
Significantly less than half of all organizations that use or develop AI have any formal or substantial governance structures for AI.”

The organizations that do have some governance structures in almost all cases are past the stage of building general frameworks for AI governance, but have not yet developed criteria or metrics to assess their effectiveness (Fig. 1). In other words, organizations that have built AI governance structures generally are at the beginning of the implementation stage.



Figure 1. Illustration of AI Governance Maturity

Among those organizations that do have some governance structures, there are some key trends with regard to implementation strategies and challenges. The key takeaways about governance structures are:

- 01 **Importance of repositories and inventories.** Organizations need to have a central inventory of AI, ML, and ADS applications to be able to govern meaningfully. This is especially the case with large organizations that often need not only inventories, but also centralized repositories with systems for collecting data about algorithms they are using for purposes of compliance, risk assessments, and risk mitigation.
- 02 **Importance of risk assessments.** All organizations we interviewed either had a clear process for risk assessments for their AI, or were in the process of developing strategies and tools for such assessments. AI risk assessments are crucial for both compliance and risk management.
- 03 **Difficulties finding employees with the right skills.** Given how new the AI ethics and AI governance field is, it is unsurprising that most organizations have struggled with assembling a team that has sufficient data science/ML skills as well as an ethics/compliance background. This was especially true for small and medium enterprises (SMEs). Many interviewees articulated uncertainty about whether data scientists should develop ethics skills, or vice versa.
- 04 **Significant difference in the maturity of AI governance between SMEs and large enterprises.** This is partly because larger enterprises tend to have a larger number of algorithms and thus face an increased risk, compared to an SME. This is also due to the fact that larger enterprises have an easier time overcoming the human resource problem mentioned in 3.
- 05 **Lack of external stakeholder engagement.** Almost across the board, organizations are not engaging with external stakeholders with respect to AI ethics and governance issues, including for purposes of risk assessments. While there is a broad recognition that there should be engagement with external stakeholders such as end-users, patients, clients, and others in the process of AI governance, most institutions do not have mechanisms to do so.
- 06 **AI governance is seen as a compliance issue.** Interestingly, despite the relative dearth of regulations regarding AI, AI governance is, in more than half of the organizations we spoke with, seen as a compliance rather than an ethics issue. This means that the most common approach to governance focuses less on holistic attempts to limit harm and more on regulatory requirements or particular issues likely to be the subject of regulation, such as bias, transparency, and privacy. This also means that organizations are less commonly embracing soft-control measures like responsible innovation standards for AI systems.
- 07 **Wide range of reasons for adopting AI governance structures.** There was not a clear pattern of motivation for building AI governance structures and tools. Some organizations were driven by upcoming regulation, others by reputational risk or media scrutiny, and some were driven by internal pressures from their own ethically-minded data scientists or legal and compliance officers.

- 08 **Culture is central to the uptake of AI governance initiatives.** Given how nascent the field of AI governance is for most organizations, there was not yet a lot of internal knowledge about what structures, initiatives, and processes are working or why. There was, however, a widespread acknowledgment that when things worked well it was because of a general culture that encouraged speaking up and taking into account ethical considerations in business decisions. At least one organization identified an increase in the willingness of people to speak up about issues relating to AI as a key potential metric to assess their AI governance system.
- 09 **AI governance sits with a range of functions.** Even though AI governance is often approached through the lens of compliance, roles and responsibilities related to governance and risk management of AI are located in different units across organizations. Sometimes the legal department is responsible for AI governance, sometimes traditional compliance officers assume this role, sometimes developers, and in still other cases personnel in business functions.
- 10 **No clear metrics emerging.** None of the organizations that we have interviewed are yet measuring the effectiveness of their AI governance initiatives. Several interviewees mentioned what they think successful governance will achieve, but there is a significant need for metrics and criteria to assess effectiveness.

The organizations we interviewed were those that are cognizant of AI harms and risks and have made some attempts to develop AI governance tools. These organizations appear to be converging on a few key approaches to AI governance. This is important because, while it has become evident in the last few years that AI ethics frameworks appear to be converging, our analysis of AI governance indicates that something similar may be occurring in the sphere of implementation, which could pave the way for preliminary AI governance best practices.