

White Paper

The AI Governance **Gap:** What Leaders and Boards Are Missing

92% of Nonprofits Are Using AI.
Most Have No Governance in Place.



2026

In early 2024, a finance employee at Arup participated in what seemed like a routine video call. His CFO and several familiar colleagues were present. The conversation unfolded normally, with clear instructions being communicated. [By the end of the call, the employee authorized transfers totaling approximately \\$25 million.](#)

\$25M

Lost by Arup after a deepfake video call bypassed human authorization entirely.

It was only afterward, upon following up with the head office, that he learned every participant on that call had been an AI-generated deepfake. No system had been breached. No firewall touched. AI was weaponized to manufacture trust and bypass human judgment.

In a separate instance that same year, [a government-commissioned report conducted by Deloitte Australia came under scrutiny after revealing multiple errors](#), including incorrect references and citations to non-existent sources, one of which was a quote falsely attributed to a sitting Federal Court Judge. These inaccuracies stemmed from undisclosed use of generative AI in the report's preparation. Deloitte later issued a corrected version and agreed to a partial refund, prompting the Australian government to signal that future contracts will require stricter AI disclosure obligations.


Though these examples may seem extreme, they highlight a concerning trend not limited to isolated incidents or large organizations. One case involved AI being used against an organization. The other involved AI being used within one. Both cases draw attention to a shared issue: Artificial Intelligence is influencing decisions and outcomes in organizations in ways that are often invisible, misunderstood, or poorly governed.

AI has likely become part of your organization's operations, regardless of whether it has been formally adopted. It rarely makes its entrance through a single, deliberate decision. Instead, it seeps in gradually. A software update introduces a new smart feature. A manager experiments with a transcription tool during a team call. A consultant uses an online model to process survey data. A project officer uploads data into a complimentary AI analysis tool. A grant management system begins automatically scoring applications. Individually, these moments seem minor, but collectively, they significantly shape how information is generated, what gets shared, how it's interpreted, and ultimately, how decisions get made.

The numbers tell a striking story. [The 2026 Virtuous report on nonprofit AI adoption](#) shows that 92% of nonprofits are using AI in some capacity. [Yet, according to a TechSoup study](#), 76% of nonprofits do not have an AI strategy, and 80% lack even a basic acceptable use policy. At the same time, 85% report strong interest in expanding their use of generative AI and predictive analytics.




of nonprofits
are using AI



have no AI
strategy in
place



lack even a
basic use
policy



report strong
interest in
expanding
generative AI

This illustrates a growing governance gap. Organizations are adopting AI at a pace much faster than they are establishing the necessary structures to oversee it, and this gap continues to widen. The same Virtuous report suggests that organizations that develop even basic governance frameworks will greatly outpace those that do not, with others falling further behind.

The risks associated with inadequate AI governance are not confined to high-profile failures like those at Arup or Deloitte. Even in the absence of dramatic incidents, unclear governance can lead to ethical dilemmas, expose organizations to legal and compliance issues, and alter the organization's overall risk profile.

When Adoption Outpaces Oversight

When boards or leaders think about AI adoption, they might picture a formal initiative that includes a steering committee, a defined budget, and a structured system rollout. However, the reality is usually quite different. AI typically integrates itself gradually, embedding itself into the everyday tools and processes that teams already rely on.

The journey often starts with small changes: a vendor introduces an AI-powered recommendation feature, a grant management system adds automated scoring, or a recruitment platform starts ranking candidates based on patterns in past hires. At the same time, staff begin using built-in tools to summarize reports or generate insights, while program teams leverage AI-generated research to inform where and how new initiatives are launched. Taken individually, these shifts can seem so minor that they don't warrant governance. Together, they begin to influence how information is produced, how options are framed, and ultimately, how decisions are made.

At first, the benefits are clear. Staff save time, uncover more insights, and experience a more data-driven approach to decision-making. In many respects, AI can significantly enhance how organizations operate.

But this is a "yes, and?" situation.

Yes, AI can save time, expand access to information, and enable more efficient, data-driven processes. And at the same time, it relies on data, assumptions, and filters that are not always visible to the user and can produce outputs that are incomplete, biased, or simply incorrect. It's like getting a test score without knowing how it was graded or whether the grading system itself is reliable.

The challenge lies in the fact that these dynamics often unfold before leadership has established a clear AI strategy or developed governance structures, and before boards or executive teams have had a structured conversation about AI's role in the organization. Consequently, a significant gap emerges between the areas in which AI already influences decisions and the organization's capacity to effectively oversee, question, and govern its use. This situation raises a critical governance question:

What safeguards are needed so that AI does not expose the organization to avoidable risks?

From there, a deeper set of considerations follows:

- ▶ Do current oversight structures account for AI-related risks?
- ▶ Do board members and senior leaders have enough AI literacy to provide informed oversight?
- ▶ How might AI reshape the organization's risk profile, internal controls, decision rights, and lines of accountability?
- ▶ How will the organization assess its readiness for broader integration?
- ▶ Are systems and teams prepared for AI to support or influence mission-critical functions?

Where AI Is Already Showing Up

Understanding where AI is embedded across the organization is crucial for implementing meaningful oversight. While the adoption of AI often begins informally, its influence tends to concentrate in a number of core functional areas.

1. Finance

In financial management, AI-enabled tools are already supporting a range of routine and analytical tasks. These tools can match spending across accounts, process large batches of expense records, and highlight discrepancies between budgets and actuals. Some systems forecast cash flow or burn rates using historic and real-time accounting data, while others flag potential duplicate payments or unusual spikes in specific cost categories.

Increasingly, AI-enabled tools are moving beyond processing and analysis into areas that were traditionally guided by human judgment. For instance, some systems can draft financial narratives that explain variances, justify budget shifts, or summarize performance for donors and boards. Others suggest which transactions or accounts should be prioritized for audit review, effectively guiding where scrutiny is applied. In certain cases, AI tools may also recommend whether an expense should be approved, escalated, or flagged for further review. Additionally, AI-generated financial risk indicators can influence how partners are perceived, shaping decisions about financial viability and risk management.

When well configured and properly monitored, AI tools can accelerate workflows and strengthen internal controls. However, when models drift from their intended purpose, are poorly calibrated, or operate without sufficient oversight, errors can compound quickly. False positives or unintended inaccuracies may compromise the reliability of financial reporting, weaken internal controls, and create exposure to regulatory and donor risk.

Finance governance questions to ask:

- ▶ What financial decisions are currently being informed or influenced by AI, and how material are they to the organization?
- ▶ Do finance, risk, and internal audit teams have sufficient understanding of how these tools operate and their limitations?
- ▶ Is appropriate human oversight maintained, particularly in exception handling and final decision-making?
- ▶ Is there clear visibility for stakeholders, including the board, on where and how automation is used in financial processes?
- ▶ Are AI models subject to ongoing testing, monitoring, and periodic review to detect drift or unintended changes in behavior?

AI Gone Rogue:

The consequences of misplaced trust in automated systems can be both immediate and irreversible. [A notable example involves the founder of SaaStr, who utilized Replit's AI coding assistant to develop a production application.](#) Despite explicitly instructing the assistant not to modify any code on 11 separate occasions, it disregarded those instructions, fabricated test data, and ultimately deleted a live production database. This incident highlights a critical flaw: the AI model had extensive access without contextual checks, approval steps, or limits on its actions, including what data it could manipulate and when. By the time the error became apparent, the damage had already been done. In response, Replit implemented basic development and production separation to mitigate such risks.

11

Instructions ignored before the production database was deleted

This cautionary tale holds significant implications for finance. An AI tool granted broad access to financial systems without adequate human oversight at crucial decision points can execute major decisions more swiftly than human operators can respond.

Issues such as a misconfigured expense approval tool, an automated reconciliation system operating beyond its intended parameters, or a fraud detection model acting on unreliable data can accumulate errors over reporting periods, often remaining undetected until significant damage has occurred.

2. Human Resources

As in finance, AI is increasingly influencing decisions that were traditionally guided by human judgment and, in this case, decisions about people, such as who gets hired, promoted, or flagged as underperforming in their roles. Recruitment platforms leverage AI to screen thousands of CVs in seconds, filtering candidates based on keywords and patterns drawn from past hires. Similarly, AI-enabled video interview tools assess candidates based on speech patterns, tone, or facial expressions.

The influence of AI doesn't stop at recruitment. It extends into the ongoing management of employees, where AI-enabled systems monitor various metrics such as work hours, email activity, and other behavioral indicators to generate performance scores. Some AI tools are sophisticated enough to predict which candidates are most likely to excel in specific roles, identify employees at risk of burnout, and recommend individuals for promotion or retention. Consequently, AI not only determines who joins an organization but also influences who advances within it.

While the efficiency gains from these technologies are significant, they come with notable risks. Screening tools may overlook talented candidates whose career trajectories diverge from conventional paths, such as those with non-traditional work experiences, or those whose names, locations or CV formats do not match historical hiring data. Video-based evaluation systems risk reinforcing cultural, linguistic, or physiological biases, as they may inadvertently profile individuals based on their appearance or mannerisms. Monitoring technologies can sometimes verge on invasive surveillance, infringing upon employee dignity and privacy. Further, the collection of extensive personal data by hiring and performance systems raises concerns about data protection and confidentiality.

3. Program Operations

AI's influence is increasingly evident in program operations, where AI-enabled tools and systems play a crucial role in determining essential programmatic factors, such as where to focus initiatives and who should receive support. For instance, advanced models can process vast amounts of data in a fraction of the time it would take traditional methods. This capability allows organizations to identify regions facing the highest risks, determine where to expand programs, and prioritize communities that require urgent assistance during emergency responses. Consequently, these systems not only shape the direction of programs but also influence the distribution of their impact.

However, the reliability of AI-generated outputs is heavily dependent on the quality of the underlying data. Many systems rely on historical or third-party data that may be incomplete, uneven, or biased. As a result, communities with weaker digital footprints can be entirely overlooked in decision-making processes. If these recommendations are not carefully reviewed, they risk perpetuating existing patterns of neglect, sidelining areas that have historically received less attention or investment. Often, the communities that are least represented in the data are also the ones in greatest need. The automated selection processes involved can also lack transparency, leaving communities uncertain about why they were included, excluded, or deprioritized.

Since AI-generated outputs are presented as data-driven, they may appear objective despite the limited underlying assumptions.

HR governance questions to ask:

- ▶ Have AI tools used in HR been reviewed for compliance with labor laws, anti-discrimination standards, privacy requirements, and data protection regulations?
- ▶ Do staff and applicants understand how their data is collected and used, and is informed voluntary consent obtained where required?
- ▶ Are there clear and accessible mechanisms for individuals to question, challenge, or seek review of decisions influenced by AI?
- ▶ Are there clear ethical guidelines and due diligence processes governing the use of AI in recruitment, performance evaluation, and staff monitoring?
- ▶ Are AI-influenced decisions documented in a way that allows for audit and explanation if challenged?

Program governance questions to ask:

- ▶ To what extent is the organization relying on AI outputs for program decisions?
- ▶ Could AI unintentionally shift program priorities away from the organization's mission or values?
- ▶ How well do AI-driven recommendations reflect local context, and where might they oversimplify complex realities on the ground?
- ▶ Are there safeguards to prevent AI from reinforcing existing patterns of underinvestment or neglect?

4. Strategic and Risk Decision Support

AI-generated dashboards and predictive models have evolved beyond their traditional role in back-office analysis; they have made their way into leadership discussions and boardroom deliberations, influencing the conversations that determine organizational direction.

For example, a foundation might use predictive models to identify which countries offer the highest potential for impact. A nonprofit may rely on AI-generated dashboards to decide which programs to scale back or expand. At the same time, fraud detection tools can flag unusual expenditure patterns or emerging compliance risks when spending falls outside historical norms.

However, the role of these tools extends far beyond mere data analysis. They actively influence what information is highlighted, which risks are deemed most urgent, and which strategies are considered viable. Over time, they begin to shape not only individual decisions, but how those decisions are framed in terms of what leadership sees, what they focus on, and what they consider possible.

Strategy and risk governance questions to ask:

- ▶ How is AI shaping the information and options that reach leadership and the board?
- ▶ Do decision-makers understand how AI-generated insights are produced, and what assumptions, indicators, or data inputs they rely on?
- ▶ Are there processes in place to critically challenge AI-generated insights before they inform high-stakes decisions?
- ▶ How is the organization balancing AI-informed recommendations with human judgment?
- ▶ How is the organization ensuring that AI-informed decisions remain aligned with its mission, values, and long-term priorities?

The challenge is that AI-generated insights are presented as data-driven, giving them an appearance of objectivity. This veneer of objectivity can be misleading, particularly when the underlying logic is incomplete or poorly understood. For example, a dashboard that appears authoritative may be built on assumptions that have not been examined, and a risk score that feels precise may reflect a narrow or outdated view of what risk looks like in a specific context.

As a result, when the underlying data is incomplete, unrepresentative, or erroneous, decisions can drift away from an organization's values and priorities. This divergence can be subtle and challenging to detect, especially since the outputs generated by these tools often appear credible and well-evidenced.

5. Informal and Invisible AI Use

Alongside more visible applications, a significant portion of AI use exists outside formal systems and approved initiatives. Not all AI operates within core platforms or structured projects; in fact, some of the most substantial risks arise from tools that have not been formally sanctioned, often because no one thought to inquire about their usage.

Many AI tools are free, easily accessible, or subtly integrated into platforms that teams already rely on for their day-to-day work. This ease of access allows them to be adopted without undergoing the typical processes of procurement, legal, or IT reviews. This phenomenon is increasingly referred to as shadow AI: tools that engage with organizational data while functioning outside formal oversight. Because these tools circumvent the very processes designed to manage risk, such as procurement checks, IT security reviews, legal sign-offs, and data governance controls, they introduce blind spots that can be challenging to detect until a problem arises.

In practice, informal AI use typically enters organizations through two channels:

1. Everyday AI in Existing Systems

The first and most common channel through which AI enters organizations is via general office tools and systems that employees use in their daily workflows. Staff members often leverage AI to draft documents, summarize meetings, prepare reports, or generate creative ideas.

- For instance, transcription tools can join video calls and produce concise summaries, while widgets can compress lengthy chat threads into brief summaries, generate task lists from emails, or autofill compliance checklists. Additionally, some applications have the capability to scrape content from shared drives and seamlessly insert it into templates with minimal human oversight.
- Because these tools are embedded within existing frameworks and are user-friendly, they are frequently adopted without thorough scrutiny. Many of them request broad access to internal platforms, often without users fully understanding the permissions they are granting. This can lead to a range of risks that are varied and sometimes subtle.

For example, sensitive or proprietary information may inadvertently be uploaded to open-source platforms. AI-generated summaries, though convincing, may contain inaccuracies that can lead to mistakes in official documents. Meeting transcription tools might record conversations without the explicit and informed consent of all participants. There's also the potential for staff to draft internal strategies or reports using AI tools, which could unintentionally expose intellectual property to services that retain or reuse submitted content.

Over time, a heavy reliance on these AI tools can erode the organization's unique voice and local nuances in communications.

To address these concerns, it is crucial to ask the right questions:

- ▶ Does the organization have clear guidelines on acceptable AI use by staff, and are staff adequately trained?
- ▶ Is data entered into third-party tools subject to privacy, donor, and contractual restrictions?
- ▶ Is there visibility into plug-ins, extensions, and applications that have access to organizational data?
- ▶ Is the organization aware of embedded AI functionalities within the systems it depends on?

2. Vendor and Partner AI Use

Perhaps the most opaque form of AI exposure comes from third parties. Vendors, consultants, and implementing partners may use AI as part of delivering contracted work or carrying out program implementation. In practice, this can take many forms. A consultant may use AI to analyze stakeholder interviews and generate findings for a report. A vendor may rely on automated tools to screen candidates or process large datasets. An implementing partner may use AI to prioritize beneficiaries, conduct surveys, or synthesize program data for reporting. AI may also be used to draft proposals, prepare donor reports, or generate monitoring and evaluation outputs.

In many instances, the organization funding or contracting the work may not be aware that AI was involved. Yet, accountability for the outcomes does not fully transfer, and the contracting organization remains responsible for the quality, integrity, and impact of the work.

The risk here is that AI's influence on outputs and decision-making can happen without visibility. For instance, errors or bias introduced during analysis may carry through into program design or strategic choices. Automated selection processes may reflect historical patterns of exclusion.

Reports generated through AI tools may rely on incomplete or flawed synthesis, while underlying data may be retained or reused by external systems, potentially exposing sensitive organizational or beneficiary information to third-party platforms without clear safeguards.

To mitigate these risks, organizations should consider several key governance questions when engaging with vendors and partners:

- ▶ Do contracts and procurement processes require vendors to disclose any AI tools used in service delivery?
- ▶ Are outputs generated with the assistance of AI subject to the same review standards as traditional deliverables?
- ▶ Are data ownership, AI ethical use, legal, and privacy protections clearly defined and enforced in third-party engagements?
- ▶ Are there restrictions in place on what organizational data vendors/third parties may feed into AI systems?

Five Ways to Start Closing the Governance Gap

The focus now shifts to what this means in practice.

Closing the governance gap does not require becoming an AI expert. It requires making AI use visible, setting clear boundaries, equipping people to ask better questions, and ensuring that accountability remains with those making decisions.

To assist with this process, here are some practical starting points:

1. Establish an AI Use Policy

The foundation for consistent and responsible use of AI within an organization is an AI Use Policy. This policy should define where AI applications are permitted, where they are restricted, and how sensitive information must be managed, both internally and in interactions with third parties. The objective is to move from a general understanding of AI use to clear expectations that staff and partners can consistently apply in practice.

2. Map AI Use, Risks, and Controls

As highlighted in earlier sections, AI is often already embedded across systems and tools already being used in the organization. An AI mapping exercise helps make this visible. This involves identifying the areas where AI is currently used, outlining where it interacts with personal or sensitive data, determining how it influences significant decisions, and recognizing where overreliance on automated outputs may occur. Mapping also highlights where failures could result in operational, financial, legal, or reputational risks.

➔ The goal is to give leadership a clear enough picture to ask informed questions and identify areas that need immediate attention.

3. Build AI Literacy Across the Organization

It is essential for everyone within the organization, including board members, senior leadership, and staff, to receive training on AI governance. However, the training should look different at each level.

Boards and executive teams require enough familiarity with AI to ask informed questions and provide oversight. Staff need clear, day-to-day guidance on how to use tools responsibly, when to question outputs, how to handle sensitive data, and when to escalate concerns. Importantly, this learning should be hands-on. Direct engagement with AI tools builds far more practical understanding than theoretical or high-level briefings alone.

4. Integrate AI into Governance and Oversight Structures

AI oversight does not need to sit within a separate committee or stand-alone framework. An effective and durable approach is to integrate it into existing governance structures.

For example:

- Risk and audit committees can request periodic reporting on AI use and associated risks, including key systems, vendors, and any incidents or near misses.
- Legal and compliance teams can monitor evolving AI laws across relevant jurisdictions and translate these into contracts, policies, and internal controls.
- Procurement processes can be updated to require disclosure of AI use, data handling practices, and responsibilities in the event of harm, error, or breach.
- IT and technology teams can maintain visibility over AI tools and integrations across the organization, ensuring that AI-related risks are reflected within broader technology governance and data protection frameworks.

For most organizations, embedding AI oversight into existing structures is both sufficient and a practical starting point.

However, for organizations adopting AI at scale or embedding it into mission-critical processes, this approach alone may not be enough. In such cases, the complexity, volume, and stakes involved necessitate a more intentional strategy. This may involve establishing a dedicated AI governance structure with clear ownership, defined accountability, formal AI risk assessment processes, and periodic independent reviews.

➔ Organizations in this latter category might consider engaging an AI governance specialist to support policy design, risk assessment, governance structures, and training.

5. Keep Final Accountability Human

Across all of these steps, one principle remains constant: AI can inform decisions, but final accountability rests with human decision-makers who can justify those decisions and, when necessary, change course.

When AI systems are used to inform or influence decisions affecting individuals or organizations, it is important to maintain accountability and explainability. Leaders must be able to articulate these decisions in clear and defensible terms. This requires understanding the criteria applied, the data relied upon, and how human judgment was used to validate, challenge, or adjust the AI-generated output.

This is particularly important in the following situations:

- Determining eligibility for funding, services, or program participation based on AI-generated assessments
- Declining or approving funding to a partner based on an AI-generated risk score
- Approving or rejecting expenditures, or classifying transactions, based on automated financial controls
- Generating financial reports, variance explanations, or donor disclosures using AI-assisted tools
- Prioritizing or excluding communities from programs based on modelled need
- Hiring, promoting, or terminating staff based on AI-assisted assessments
- Flagging fraud, compliance risks, or triggering sanctions based on automated detection systems
- Processing or analyzing personal, financial, or beneficiary data through AI systems, particularly where data is transferred, retained, or reused by third-party tools
- Making decisions where individuals or organizations may be denied access or opportunity without a clear explanation or the ability to challenge the outcome

➔ In each of these situations, decision-makers must be able to explain the reasoning behind their decisions, instead of merely referencing the system that produced them. In any legal or regulatory context, the standards do not change due to the involvement of AI. Responsibility does not transfer to the model; it stays with those who make, approve, rely on, and act upon decisions influenced by AI.

Simply citing that the system recommended it or the AI produced it does not constitute a valid defence; at best, it serves as a statement of fact.

Moving Forward

Closing the governance gap starts with clarity and understanding of where AI is already shaping decisions, and with ensuring that leadership is equipped to engage with it deliberately, not reactively.

At Humentum, we recommend beginning with a structured AI leadership session, which is a focused conversation that moves beyond general awareness to address the questions that matter most for your organization.

- ▶ Where is AI already influencing decisions, and how visible is this to leadership?
- ▶ What risks does this introduce, and where are the gaps in oversight?
- ▶ What does responsible AI use look like in the context of your mission and values?
- ▶ How should governance and accountability be structured, and who owns it?

These questions require honest internal reflection, the right people in the room, and the willingness to act on what the conversation surfaces.

Is your leadership team ready to close the gap?

Humentum works with boards and senior leadership teams through facilitated [Leadership in the AI Era sessions](#), designed to move organizations from awareness to informed action. If your leadership team is ready to take that step, we would welcome the opportunity to be part of it.

Email client@humentum.org to discuss an AI leadership session at your organization.

About the author



Shila Nhemi

Technical Director, Humentum

Shila Nhemi has over 18 years of experience across private and nonprofit sectors, specializing in corporate governance, risk management, board practices, compliance, institutional strengthening, and AI Governance.

Before joining Humentum, Shila worked with KPMG, Toyota, IFRC, and the Norwegian Red Cross and has advised organizations across Africa, Asia, the Middle East, the Americas, and Europe. She is a qualified corporate governance lawyer, Chartered Governance Professional, Certified Public Accountant, and is currently pursuing her Doctorate in Law, researching AI Regulation. Shila also holds certificates in AI Governance, AGI Strategy, and AI Policy Research.

Related publications from the author:

- [*Do No Harm: What Happens After Someone Comes Forward*](#)
- [*Beyond grants: 7 alternative funding models for nonprofits*](#)
- [*Five Initiatives for Successful Cost Recovery*](#)



www.humentum.org