Intelligent Organizations

Governing AI Under Uncertainty: What Boards Can Learn from Global Cases

National Director Conference 2026

Kathleen D Kennedy
Executive Director, MIT Center for Collective Intelligence
Senior Director, MIT Open Learning
kdkenn@mit.edu

Center for Collective Intelligence



AI and emerging technologies are changing the world.
We see new things every day.



The New York Stock Exchange transformed so trading will be extended to 22 hours per day

10.22 stock value: ▲ $26.31

11.22 stock value: ▼ $1.19

© 2026 MIT







TESLA ? Alphabet









Volatility in markets

© 2026 MIT



Completely new modes of operating are emerging

- Downtown to JFK in 7 mins

© 2026 MIT

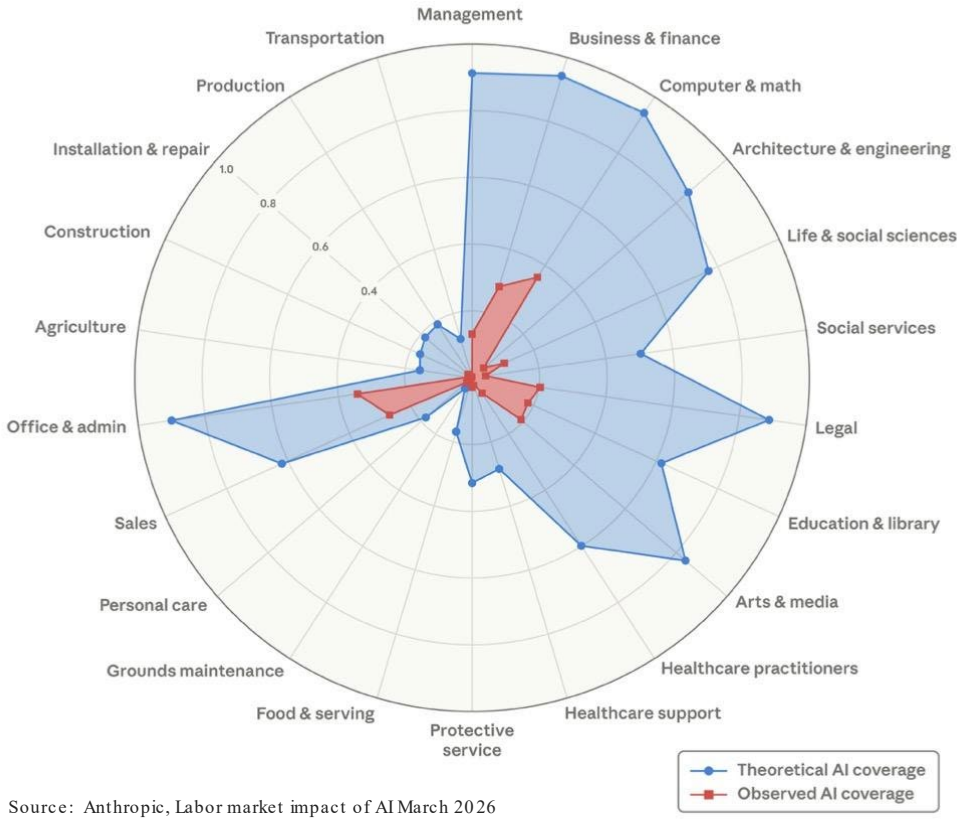


AI models can detect cancer, often better than humans.

Source: MIT News

© 2026 MIT

AI is reshaping work

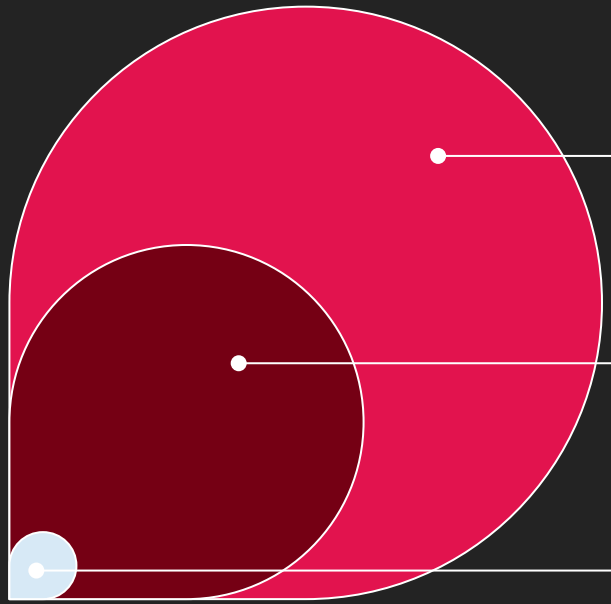


Source: Anthropic, Labor market impact of AI March 2026

Figure 2: Theoretical capability and observed exposure by occupational category
This figure shows the share of job tasks that LLMs could theoretically perform (blue area) and our own job coverage measure derived from usage data (red area).

© 2026 MIT

Organizations continue to invest in AI transformation…



60% of companies plan to increase their transformation investment in 2025

36% Will maintain their 2025 investment level

04% Wil reduce their digital transformation investment in 2025

Statistics from: BCG X, Mind the Tech Gap

© 2026 MIT

but they are finding it hard to transform…

95%

of organizations got zero return on their investment in Generative AI pilots

State of AI in Business 2025, MIT NANDA (Networked Agents and Decentralized Architecture)

© 2026 MIT

How should our organization and governance adapt now that AI is changing work?

© 2026 MIT

Intelligent Organizations

© 2026 MIT

Intelligence

=

?

© 2026 MIT

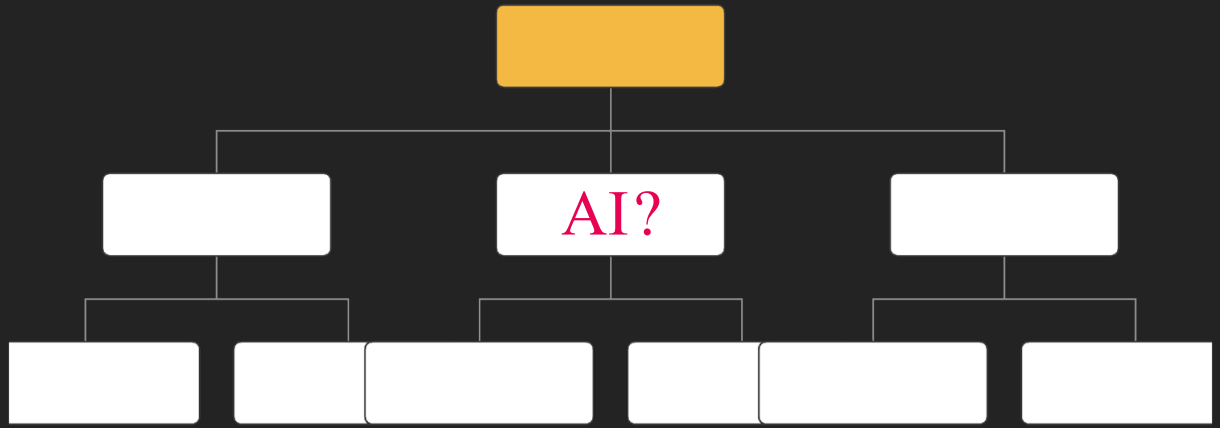
Intelligence

=

The ability to achieve goals

© 2026 MIT

Where is AI in the organization?



AI has been in many organization for a while

But the roles are changing quickly…



Tools

Execute



Assistants

Support

© 2026 MIT

AI Is Taking on New Organizational Roles

Increasing Autonomy → *Increasing Organizational Influence* → *Increasing Governance Responsibility*



Tools

Execute



Peers

Collaborate



Assistants

Support



Managers

Direct

© 2026 MIT

How work is done and decisions are made is changing

Roles

AI is changing how work gets done.

- Which activities are now performed by AI?
- Which still require human judgment?
- Which jobs are being redesigned?
- What new capabilities are required?

Objective & Incentives

Humans follow incentives
AI only has objectives.

- What objective has the AI been given?
- Does it optimize the right outcomes?
- Are human incentives consistent?
- Could misalignment create unintended behavior?

Authority & Accountability

Who decides?
Who is responsible?

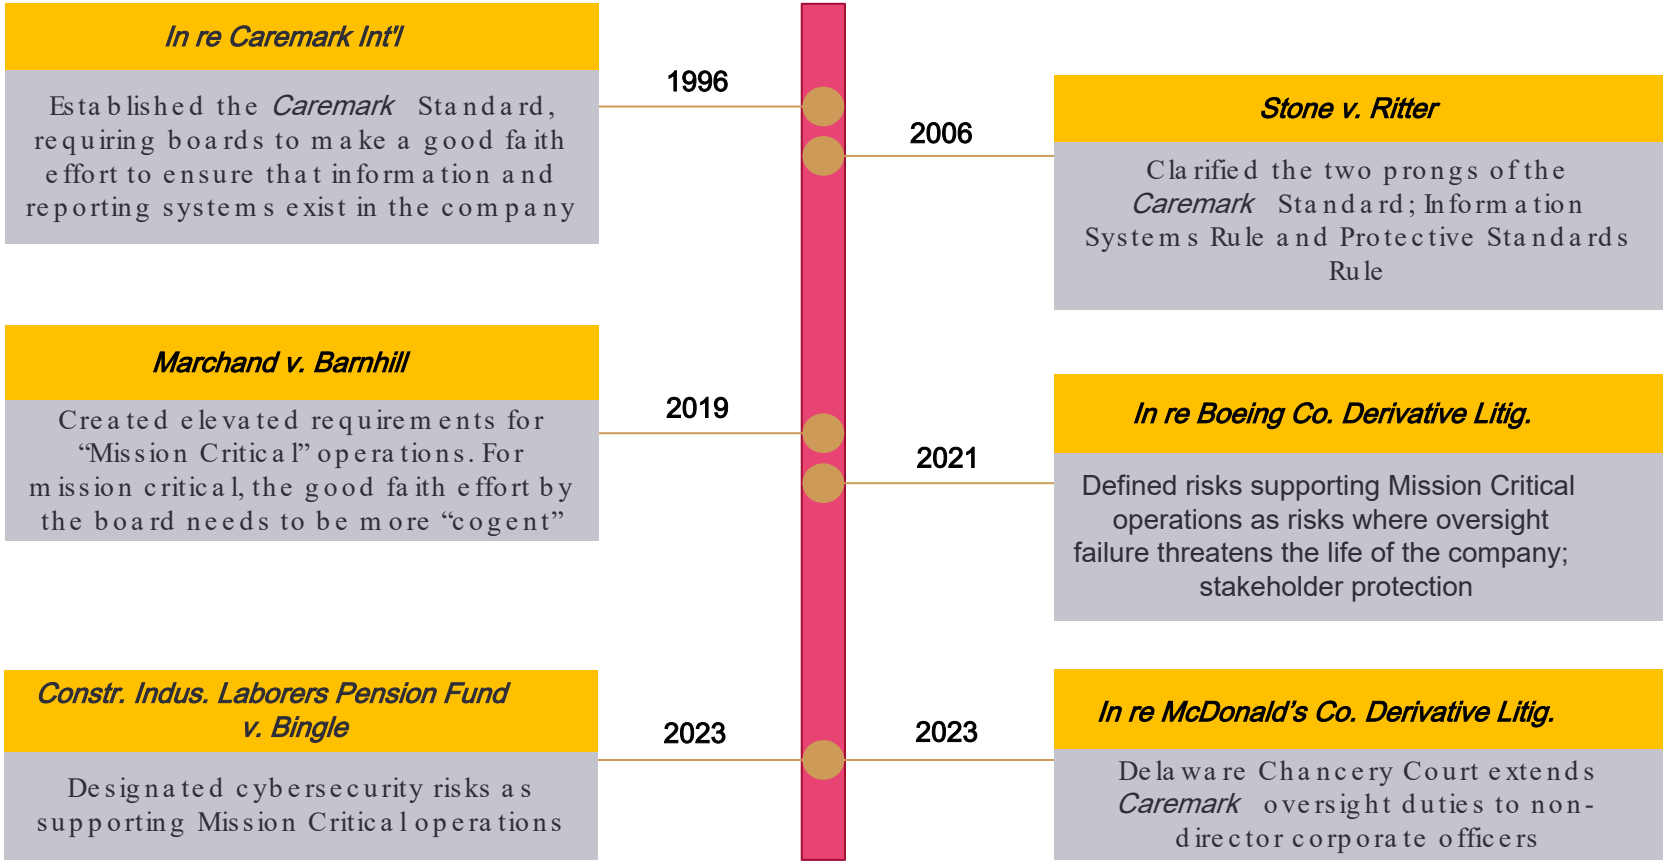
- What decisions has AI been given?
- What authority remains with humans?
- Who is accountable when AI contributes?
- Can humans intervene when necessary?

How does governance need to change as AI transforms organizations?

Global Cases

© 2026 MIT

Cases that establish minimum oversight standards



© 2026 MIT

Three Lenses

Roles

AI is changing how work gets done.

- Which activities are now performed by AI?
- Which still require human judgment?
- Which jobs are being redesigned?
- What new capabilities are required?

Objective & Incentives

Humans follow incentives
AI only has objectives.

- What objective has the AI been given?
- Does it optimize the right outcomes?
- Are human incentives consistent?
- Could misalignment create unintended behavior?

Authority & Accountability

Who decides?
Who is responsible?

- What decisions has AI been given?
- What authority remains with humans?
- Who is accountable when AI contributes?
- Can humans intervene when necessary?

Deloitte Australia

What Happened?

- Government report prepared using Gen AI, but Deloitte remained the author.
- Fabricated citations remained in the final document.
- The report was withdrawn.
- *Key lesson: AI-generated work still requires human governance.*

© 2026 MIT

Deloitte Australia

Roles

- AI used to produce a report, but Deloitte remained expert adviser and author.
- AI moved beyond administrative support into research, analysis, citation, and drafting.
- Humans were expected to verify the work, but fabricated sources and quotations passed through the review process.
- The case shows that AI may perform parts of the expert role, while the human and the firm retain the professional role.

Objective & Incentives

- AI offered faster research and drafting, reducing the time and cost
- Deloitte was incentivized to capture those efficiency gains while delivering work that appeared authoritative and complete
- Reviewers may have been inclined to trust polished AI-generated material, particularly when citations looked credible.

Authority & Accountability

- The AI had authority to shape the report's content, but no responsibility for whether that content was true.
- Deloitte remained accountable because it accepted the commission, exercised professional judgment, and delivered the report under its name.
- The partial refund reinforced a basic governance principle: **using AI does not transfer accountability away from the organization using it.**

SCHUFA, Germany

What Happened?

- AI-driven credit scoring significantly influenced lending decisions.
- Individuals had limited ability to understand or challenge the system.
- European regulators emphasized the need for meaningful human oversight.

© 2025 MIT

SCHUFA, Germany

Roles

- SCHUFA appeared to be information provider and the bank appeared to be the decision-maker
- The case shows that an AI system can function as the **real decision-maker** even when a human formally makes the final decision.
- Governance must examine **actual influence**, not simply assigned job titles or formal roles.

Objective & Incentives

- SCHUFA's system was optimized to predict repayment risk efficiently and consistently.
- Banks incentivized to rely heavily on the score; it **reduced cost, time, risk.**
- Neither party had a strong incentive to question or explain an apparently authoritative score.
- Efficiency and risk reduction vs **fairness, transparency, individual context, ability to contest decision.**

Authority & Accountability

- SCHUFA only produced information; the bank made the lending decision.
- The bank relied on an external expert scoring system.
- This created an accountability gap: **SCHUFA exercised substantial authority without accepting responsibility for the final outcome.**
- The court effectively linked accountability to consequential influence: when an automated score largely determines an outcome, its provider cannot treat it as merely advisory.

Air Canada

What Happened?

- An AI chatbot gave incorrect information about bereavement fares.
- A customer relied on the advice.
- The court held Air Canada responsible.

© 2025 MIT

Air Canada

Roles

- Chatbot served as a customer-service representative, interpreting Air Canada's bereavement-fare policy, giving incorrect info
- Air Canada treated the chatbot as merely a communication tool, but the customer reasonably experienced it as an authorized representative of the airline.
- The case shows that when AI performs an organizational role, customers may reasonably rely on it as they would a human employee.

Objective & Incentives

- Chatbot operational objective was to answer questions and resolve customer needs without human intervention.
- Air Canada captured the efficiency benefits, while the customer bore the risk that the information might be inaccurate.
- The case illustrates the danger of optimizing AI for responsiveness and cost reduction without giving equal weight to accuracy, policy compliance and escalation when uncertainty exists.

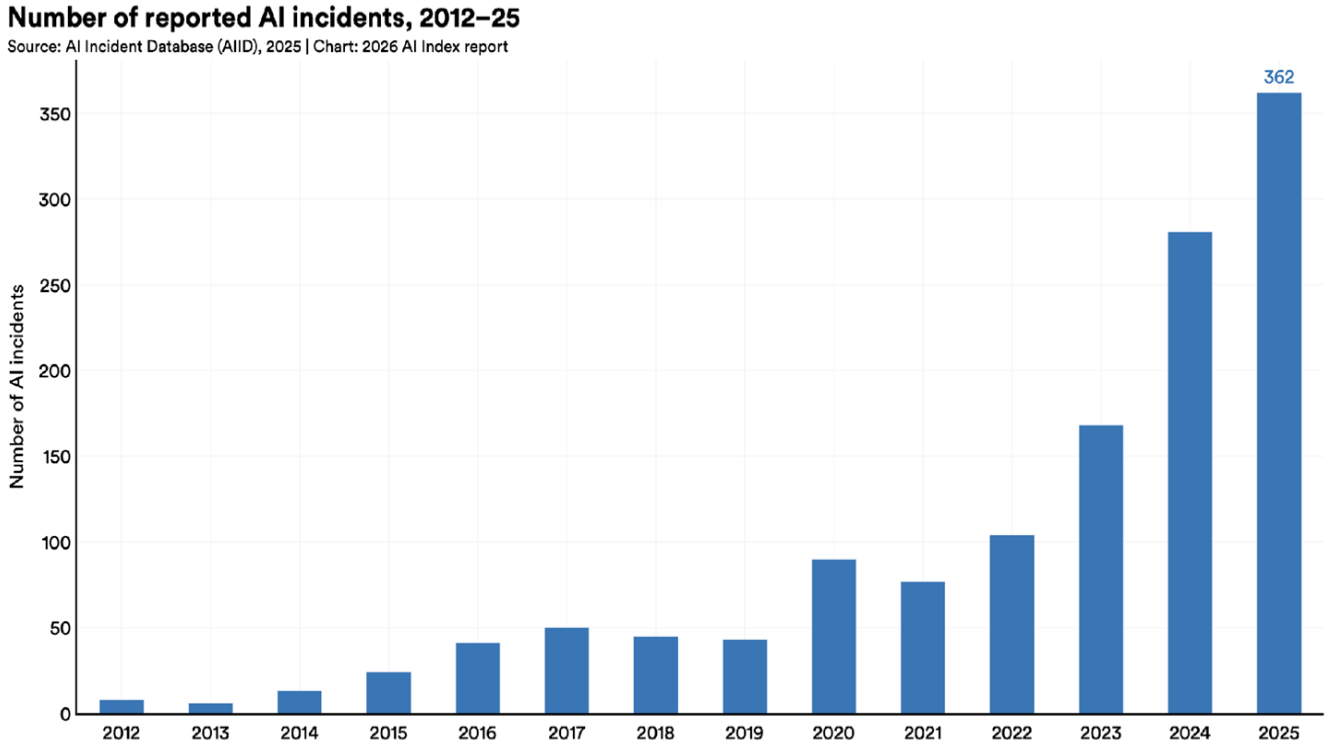
Authority & Accountability

- Air Canada argued that it should not be responsible for the chatbot's inaccurate statement, pointing to correct information elsewhere on its website.
- The tribunal ruled: the chatbot was part of Air Canada's website, and responsible for the information it provided.
- An organization cannot give an AI system authority to represent it while disclaiming responsibility when the system is wrong.
- **Authority may be delegated to AI, but accountability remains with the organization.**

What about Agentic?

© 2026 MIT

Reported AI Incidents



Source: Stanford HAI: 2026 AI Index Report

Agentic AI: the control problem

Why existing AI governance frameworks fall short

57%

of orgs deploy agents for multi-stage workflows



Principles Gap

Fairness & transparency describe values, not who may authorize, monitor, or revoke an agent



The Tool Era

Agents interpret goals & act across systems; governing them as passive tools leaves authority undefined



Accountability Gap

No lifecycle controls = no traceable ownership; when failures occur, responsibility cannot be assigned

The Key Distinction

Capability

What the system can technically do

Access

What systems or data it can reach

Authority
What the organization permits it to do; this is what must be governed

The Reframe: AI agents as licensed organizational actors

Governance question shifts: not "who has access?" → "what is this agent authorized to do?"

Licensed Agent

Authority is explicit:

Granted | Bounded | Monitored | Owned | Revocable

Each attribute is a control point. Authority does not flow from access; it is issued, scoped, and can be withdrawn (or expanded).

Human-Actor Analogy

Not because agents are people, but because the control logic transfers:

- Role definition
- Performance review
- Termination
- Supervision
- Escalation
- 'Reward' strong performance?

Two Institutional Homes

Internal Audit
Accountability — control testing, evidence, escalation, sign-off

HR
Lifecycle: onboarding, role design, performance, offboarding

Agentic AI governance: what this looks like in practice

Adopt a unified lifecycle, with governance that scales to the authority delegated

Current State – Fragmented Controls

- Committees
- Access limits
- Human sign-off
- Informal escalation

Target State – Unified Lifecycle



Bottom Line – Every Agent Needs

Identity	**Role**	**License**	**Record**	**Supervisor**	**Exit**
who it is	*what it does*	*what it may do*	*what it did*	*who owns it*	*how it ends*

How to adapt?
Is Data Governance a Good Model for AI Governance?

© 2025 MIT

Data Governance

Managing your organization's data assets

Data governance is the framework of policies, standards, and accountability structures that ensure data is accurate, accessible, consistent, and used appropriately across an organization.

Data Quality

- ✓ Accuracy, completeness, and consistency of data records;
- ✓ ensuring your data can be trusted by decision-makers.

Data Access & Privacy

- ✓ Who has access to what data, under what conditions;
- ✓ compliance with GDPR, CCPA, HIPAA, and sector-specific regulations.

Data Stewardship

- ✓ Ownership, lineage, and lifecycle management;
- ✓ tracking where data originates and who is accountable for it.

AI Governance

Managing how AI systems behave

AI governance is the set of policies, processes, and oversight mechanisms that ensure AI systems operate safely, fairly, and transparently, and align with organizational values and legal obligations.

Model Risk & Safety

- ✓ Validating that AI outputs are reliable, unbiased, and don't cause unintended harm;
- ✓ especially in high-stakes decisions.

Transparency & Explainability

- ✓ Ensuring AI objectives are understood; and aligned
- ✓ decisions can be audited, and explained to regulators, employees, customers, and other key stakeholders.

Accountability & Ethics

- ✓ Defining who is responsible when AI fails and embedding ethical principles into system design and deployment.

How They Differ

Dimension	Data Governance	AI Governance
Primary Focus	Data assets: quality, privacy, and access	AI systems: behavior, decisions, and impact
Core Question	"Is our data trustworthy and compliant?"	"Is our AI safe, fair, and accountable?"
Key Risks	Inaccurate data, breaches, regulatory fines	Biased outputs, opaque decisions, reputational harm
Role/Authority	Chief Data Officer, Data Stewards	AI Governance Board, Model Risk, Legal, CAIO, Chief Compliance/Ethics Officer
Maturity	Established practice (10 –15+ years)	Rapidly evolving (regulatory landscape in flux)

The Bottom Line

You need both, but they're not interchangeable

1. **Data Governance is the Foundation**

   Without trusted data, AI models inherit the problem ingested into the AI system. Strong data governance is a prerequisite for responsible AI programs.

2. **AI Governance is the Next Frontier**

   Even with perfect data, AI systems can still behave unexpectedly. AI governance addresses what happens inside the model, not just what goes in.

3. **Build One Unified Framework**

   Leading organizations connect both disciplines; with shared ownership, common vocabulary, and integrated risk management across data and AI.



“Increasingly, success will come, not from having more productive organizations, but from having more intelligent ones.”

– Thomas W. Malone,

MIT Professor and Author of *Superminds* and *The Future of Work*

© 2023 Proprietary & Confidential

Thank you!

Governing AI Under Uncertainty: What Boards Can Learn from Global Cases
National Director Conference 2026

Kathleen D Kennedy
Executive Director, MIT Center for Collective Intelligence
Senior Director, MIT Open Learning
kdkenn@mit.edu

Center for Collective Intelligence