Reliable Leadership Advice™



Future of Work Series • Article 2

What Does AI Mean for Your Organisation—and how Should you Prepare?

Authors:



Louise W. Björk
Deputy CEO, Governance and Transformation,
AiTek FACTORY; NED MU AI Advisory Board



Richard Moore
CEO, Partner & Board Director at MU

@ richard.moore@mercuriurval.com

+44 771 193 0760

This series of articles by MU CEO Richard Moore, and expert guests, offers help in understanding technology and the future of leadership. What will AI and other frontier technologies mean for you as a leader?

- Article 1: What AI opportunities and risk lie ahead, and how will you successfully navigate them?
- Article 2: What does AI mean for your organisation – and how should you prepare?
- Article 3: What skills and knowledge do you need as a leader in the AI-era, and how will you identify, select and develop them?
- Article 4: How will AI impact decision-making, organisational capability, and the way work and judgment is distributed across your organisation?
- Article 5: Responsible human centred AI, what does it mean for you, your organisation and your customers?
- Article 6: How to lead technology driven change? (And what's special about the coming AI change)

By looking at, and beyond, the new technology agenda, our goal is to give practical advice for leaders.

What Does AI Mean for Your Organisation – and how Should you Prepare?

AI and other frontier technologies are rapidly moving from isolated experiments into everyday work, customer interactions and core business processes. For leaders, the central question is no longer whether to engage with AI. It is how to make sound choices about where it creates value, what it changes across the organisation, and what needs to be in place before it is deployed at scale?

These questions create a difficult strategic challenge. Organisations need to improve today's performance while preparing for ways of working, decision-making and customer engagement that are still emerging.

In the *first article in this series*, we explored the opportunities and risks that AI creates and the need to navigate both deliberately.

Having understood the broader opportunities and risks, organisations must decide where AI creates real strategic value, how work and customer experiences should change, which human capabilities must be retained and developed, and what needs to be prepared before new technology is deployed at scale.

Louise W. Björk brings a particular focus on AI decision-making, organisational preparation and the practical consequences of introducing AI into real workflows.

Together with Richard Moore's perspective on leadership and organisational performance, this article explores how leaders can make sound AI decisions and prepare their organisations for what follows.

How Should I Approach AI Decisions?

The conversation pointed to a clear conclusion: AI should not be approached as a technology initiative in isolation. It is a business-change decision. The starting point is not the tool, but the problem the organisation is trying to solve, the value it wants to create and the consequences it is prepared to manage.

While AI is the immediate focus here, the same decision discipline can help leaders approach other frontier technologies: Start with purpose, understand what will change, prepare deliberately, and assess the outcome.

So, how should organisations approach these decisions? What needs to be prepared across strategy, people, data, governance and customer experience? And how can leaders measure not only whether AI works technically, but whether it creates the outcomes they intended? In simple terms, follow this flow:

Purpose → Consequence → Preparation → Evaluate

For example, take an organisation using AI to help identify a treatment for a disease:

PURPOSE: Can we identify promising treatments faster, or find relationships in the data that researchers could not reasonably identify at the same scale?

CONSEQUENCE: If it works, what changes for researchers, clinicians, patients, workflows, data use and potentially access to treatment?

PREPARATION: What evidence, specialist expertise, validation, permissions, data quality and human oversight are needed before findings influence real clinical decisions?

EVALUATE: Did it actually improve discovery, speed, accuracy or patient outcomes, and is the evidence strong enough to go further?

Start With the Business Problem, not the Technology. What is Your Purpose?

The first step comes before AI. Leaders need to establish what problem they are trying to solve, what opportunity they want to pursue, or what result they need to achieve. Only then should they ask whether AI can help.

It is easy to begin with statements such as, "We need to bring AI into the business," or, "Our competitors are using AI." But neither is a strategy. Before selecting a tool, leaders should be able to explain the business issue clearly: What needs to improve, for whom, and why now?

That means asking:

- What outcome are we seeking – better quality, faster service, lower cost, greater innovation, improved customer experience or reduced risk?
- Who will receive the value – customers, employees, partners, shareholders or society?
- What does success look like in practical terms?
- Is AI genuinely the best intervention, or would a simpler process redesign, better data, new capability or organisational change solve the problem more effectively?
- What should remain distinctly human?

This final question is especially important. In many settings, human relationship, judgement, empathy and contextual understanding can be part of the value customer is buying. A hospitality employee may notice a customer need that a system cannot interpret. A clinician, adviser or case handler needs to interpret nuance, history, emotion or ethics in a way that makes the customer feel they have genuinely been understood. Leaders should therefore ask both "Can this be automated?" and "If we automate this, do we remove part of the value the customer actually wants from us?"

The objective should be to make deliberate choices about what technology should do and what people must continue to own.

Purpose also determines the appropriate level of ambition. A drafting tool that supports an individual employee is fundamentally different from an AI agent that acts across systems, influences customer interactions or changes how roles are designed. Once the purpose is clear, leaders can decide what kind of AI capability is required, at what scale and with what authority.

Then Consider Consequences. What, Exactly, Will Change for you and Your Customer?

AI consequences need to be considered in two ways: First, what is distinctive about AI itself; second, how those effects may travel through the organisation.

Why AI Requires Additional Discipline

Many disciplines of successful technology implementation still apply: Clarify the business case, involve users, establish governance, manage change and measure outcomes. But AI adds characteristics that can make consequences faster, less visible, easier to scale and harder to reverse.

AI raises the stakes. The faster the system can act, scale and influence decisions, the stronger the evidence, oversight and control need to be.

1. Speed and scale

AI-generated outputs, recommendations and actions can spread across teams, customers and connected systems faster than people can review them. A small error, flawed assumption or inappropriate prompt can become a large organisational problem very quickly. Leaders need to design monitoring and limits for the scale a system could eventually reach – not only for the initial pilot.

2. Autonomy and reach

Agentic AI can move beyond providing advice. It may retrieve information, use tools, initiate workflows, communicate with customers, make transactions or instruct other systems. This creates a fundamental governance question: What can the system technically do, and what has the organisation actually given it permission to do? These are not the same thing.

3. Variable and difficult-to-explain outputs

Generative systems may produce different answers to similar questions, and they can sound highly convincing when they are wrong, incomplete or inappropriate. It may also be difficult to reconstruct why a model produced a particular result or why a chain of automated actions developed as it did. Verification must therefore match the potential consequence.

4. Dependency on models, data, integrations and vendors

An AI system depends on a combination of data, prompts, models, integrations and external providers. Any of these may change over time. An approval made today does not guarantee that the system will behave in the same way tomorrow. Organisations need to understand what happens if a data source changes, a model is updated, a vendor alters its service or an integration fails.

5. Unusually broad accessibility

Formal governance may be in place for approved uses, while employees independently adopt unofficial tools, upload information or build private agents to help them work. This can reveal valuable innovation and demand, but it can also create unmanaged data, security, quality, compliance and intellectual-property risks. The official implementation may not be the full AI footprint.

There is also a human risk. As AI becomes useful and usually correct, people may gradually defer judgement to it. Over time, automation can also weaken the practical experience through which individuals and organisations build expertise. If AI takes over the work through which junior professionals learn to spot patterns, handle ambiguity, make mistakes and improve, leaders should ask a difficult question "Where will tomorrow's experts come from?" What looks like an efficiency gain today could cause a capability gap tomorrow.

As explored in *the first article* of the Future of Work series, AI's opportunities must be considered alongside its unintended consequences, including the risk of crossing ethical or regulatory boundaries. The more autonomous, far-reaching, opaque and difficult to reverse an AI application is, the stronger the case for evidence, oversight and governance before it is deployed at scale.

Follow the Consequences Before Discussing Readiness

There is no useful generic state called "AI ready". Readiness depends on what the organisation is planning to introduce and what consequences or change it will create.

Before moving ahead, leaders should trace the likely effects across the organisation. This can be done by using consequence lenses as a scan to identify which consequences are relevant to the use case in front of you, and to understand where cost, risk or dependency has moved as a result of the change:

- **Strategy:** Does the technology support the organisation's real priorities, competitive position and longer-term direction? What is difficult to reverse and what does that mean?
- **Customers:** Will it improve customer value, speed, relevance and service, or reduce trust, quality or the human experience?
- **People:** Which roles, skills, career paths and responsibilities will change? Where will tomorrow's experts come from?
- **Data:** Is the information accurate, current, relevant and appropriate for the intended use?
- **Workflow:** Where do decisions and work pass between people and systems? Where are handovers most vulnerable?
- **Governance:** Who decides, challenges, overrides and stops the system? Who remains accountable? Where does practical power move?
- **Operating model:** Do ownership, incentives and decision rights still work when AI is embedded in the workflow?

Customer value should be treated as a design question, not just as a measure at the end. Leaders need to decide where automation improves speed, convenience or relevance, and where it reduces trust, empathy or the quality of the relationship. An efficient AI-enabled service can still fail if customers feel misunderstood, excluded or unable to reach a person when the situation requires judgement. In some contexts, the human interaction is not an expensive layer around the offering; it is the offering.

Human oversight also needs to be designed carefully. The idea of a single "human in the loop" can be misleading. Oversight is not meaningful simply because someone is technically present. If people review large volumes of mostly accurate AI-generated outputs, they could gradually stop questioning the system. Falling override rates may indicate better system performance, but it can also show that people have become over-reliant on it.

Effective preparation therefore goes beyond writing a policy. It may mean redesigning the workflow, limiting permissions, improving data, adding escalation points, investing in learning, slowing the rollout or deciding not to automate a particular activity.

Prepare Well: Consider at Least Data, People and Users, Measurement, Permissions, Regulations, Human Oversight

Match Preparation to Reversibility

Not all AI decisions deserve the same level of evidence or governance. A useful principle is reversibility: How difficult would it be to undo this decision if the expected value does not materialise, or if unintended consequences emerge?

A limited trial of a low-risk drafting tool may be relatively easy to reverse. Delegating significant authority to an AI system, redesigning roles, removing a human control point or changing customer interactions at scale is not. These decisions can affect trust, capability, and accountability long after the original system has changed.

The less reversible the decision, the stronger the case for:

- Evidence that the problem and proposed solution are well understood.
- A clear view of customer, employee and organisational consequences.
- Proportionate governance, ownership and accountability.
- Meaningful human oversight and the ability to intervene.
- Appropriate safeguards, permissions and escalation routes.
- Pre-agreed conditions for scaling, redesigning, pausing or stopping.

This does not mean waiting for perfect certainty. It means matching the level of preparation to the potential consequence of the decision. Leaders should be able to experiment quickly where risk is contained, while applying greater discipline when the decision is difficult to reverse.

Prepare People to Work With AI, not Simply Around it

The most important people capability is adaptability, both at an individual and organisational level. Employees need to recognise what is changing, identify emerging problems, speak up when something is wrong and learn from what is working.

However, adaptability alone is not enough. Organisations also need three connected capabilities.

1. Practical AI literacy

People do not all need to become AI engineers. But they should understand enough to use AI responsibly: how data affects outputs, what the limitations are, when sensitive or client information should not be used, and what risks may arise from relying on an AI-generated answer.

2. Domain expertise and critical judgement

AI does not remove the need for professional expertise. It makes it more important. People need enough knowledge of the business, customer, profession or process to judge whether an output is credible, helpful and safe.

Critical judgement means asking: Is this correct? Is it appropriate in this situation? What evidence supports it? What has the system missed? When should we not delegate?

Introducing AI does not remove organisational accountability. Responsibility for how the system is authorised, governed and used still sits with people and the organisation.

3. Deliberate human–machine collaboration

Leaders should be clear about what to delegate, what to retain, how to validate outputs and when people need to intervene. The aim is not to encourage employees to outsource their thinking. It is to use machines for routine, repetitive or data-intensive tasks while protecting the distinctly human capabilities that create value: judgement, imagination, relationship-building, empathy and responsibility.

This has important implications for learning and role design. If AI takes over work through which junior employees have traditionally developed knowledge, judgement and context, organisations need deliberate alternatives: mentoring, supervised practice, structured learning or new development pathways. Otherwise, short-term efficiency could weaken the expertise the organisation will need in the future. The question is not just how to make experienced people more productive today; it is how to build the next generation of people who can exercise sound judgement tomorrow.

Follow up on Decisions: Evaluate AI Effectiveness

Measure Outcomes, Behaviour and Wider Consequences

Measuring AI implementation begins with a precise purpose. If the organisation has not clearly defined the problem it is trying to solve, it will struggle to judge whether the technology has helped solve it. Leaders need to be clear about what the AI is doing, for whom and in relation to which business objective.

The measures should reflect the value the organisation intended to create. This can include cost reduction, faster response times, improved quality, productivity, revenue growth, better decisions or a stronger customer experience. But measurement should not end with technical performance, system usage or one headline KPI.

AI can also affect employee experience, customer trust, capability development, error rates, creativity, resilience and the quality of communication. These effects can be more difficult to quantify, but they are central to whether the change succeeds over time. A system may appear technically successful while customer value, trust or long-term organisational capability deteriorates.

Leaders should also align different perspectives on value. A sales leader may focus on faster responses to proposals and stronger salesforce engagement, while the board may focus on reducing the cost to sell. Both can be valid measures, but they need to be connected to a shared view of the intended outcome and the trade-offs the organisation is prepared to make.

Measurement should also support decision-making, not simply reporting. From the beginning, leaders should agree what evidence would justify scaling, redesigning, pausing or stopping the intervention. This is particularly important where the system's reach, authority or impact makes the decision harder to reverse.

The measures should therefore cover:

- The intended business outcome and the baseline against which improvement will be judged.
- The behaviours expected to create that outcome, including how people use, trust, challenge or adapt around the system.
- The customer, employee and capability effects that may emerge over time.
- Quality, speed, error, cost and risk implications.
- Whether safeguards, oversight, training and workflow redesign are working in practice.
- Pre-agreed thresholds for customer impact, quality, adoption, human overrides, incidents or unacceptable cost.

This matters because a measure can become misleading once it becomes a target. People may optimise the number rather than the outcome it was intended to represent. Leaders should therefore investigate what sits behind the measure, rather than assume that a positive-looking number represents success. Human override rates should for example be interpreted alongside evidence of whether people are still exercising appropriate judgement.

AI decisions should also be formally revisited at agreed points. Models, vendors, data, regulation and organisational circumstances can all change. Measuring from the outset – and remaining open to evidence after deployment – allows leaders to learn, adjust and make a better decision about what to do next.

Leadership: Balancing Today's Business With Tomorrow's Organisation

These decisions also place new demands on leadership. AI calls for ambidextrous leadership: the ability to run today's operation effectively while preparing for a future in which work, decision-making and economic structures may be different.

Leaders do not need to become technical specialists. But CEOs, boards and executive teams need enough fluency in AI concepts, forms of AI, data, machine learning, risks and the relevant ethical and regulatory context to ask informed questions and know when specialist challenge is required.

There may also be a shift in organisational influence. Formal authority remains with the board, CEO and executive team. But practical influence may move towards those who understand, configure or control the systems through which work increasingly flows. That includes the CTO function, but it can also include architects, product owners, data leaders, vendors and operational teams close to the technology. Leaders therefore need enough understanding and enough proximity to challenge purpose, consequences, permissions, trade-offs and evidence without trying to become engineers themselves.

They also need clarity, confidence, communication and influence. AI is being introduced in an environment shaped by both hype and fear. Employees might worry about job security, status, capability or loss of control. Customers might question the authenticity, safety or humanity of the service they receive. Leaders must address those concerns directly and honestly.

For a deeper consideration of the leadership capabilities required, readers should turn to the third article in the Future of Work series.

A Practical Route to a Sound AI Decision

A strong AI conversation should finish with a decision, not simply a discussion. Leaders can use four questions to structure that decision:

1. **Purpose:** What business problem or opportunity are we addressing, and is AI the right response?
2. **Consequence:** If this works, what will change for customers, people, workflows, data, governance and the operating model? Which of those consequences matter most in this case, and where has cost, risk or dependency moved?
3. **Preparation:** What capabilities, ownership, controls, learning, data and change support must be in place?
4. **Evaluation:** What outcome will we measure, what evidence is required, and what would make us proceed, pilot, redesign, pause or stop our decision?

This gives a practical way to connect AI decisions to strategy, people capability and customer value, while preparing the organisation for both intended and unintended consequences.

The most effective organisations will not be those that adopt AI fastest for its own sake. They will be those that connect AI choices to a clear purpose, prepare seriously for their consequences, protect human judgement and customer trust, and remain willing to adapt their approach as evidence emerges.

AI and frontier technologies can create significant value. But that value will come not from the technology alone, but from the quality of the strategic and human decisions made around it.

This gives a practical way to connect AI decisions to strategy, people capability and customer value, while preparing the organisation for both intended and unintended consequences.

The most effective organisations will not be those that adopt AI fastest for its own sake. They will be those that connect AI choices to a clear purpose, prepare seriously for their consequences, protect human judgement and customer trust, and remain willing to adapt their approach as evidence emerges.

AI and frontier technologies can create significant value. But that value will come not from the technology alone, but from the quality of the strategic and human decisions made around it.

This gives a practical way to connect AI decisions to strategy, people capability and customer value, while preparing the organisation for both intended and unintended consequences.

The most effective organisations will not be those that adopt AI fastest for its own sake. They will be those that connect AI choices to a clear purpose, prepare seriously for their consequences, protect human judgement and customer trust, and remain willing to adapt their approach as evidence emerges.

AI and frontier technologies can create significant value. But that value will come not from the technology alone, but from the quality of the strategic and human decisions made around it.

To find out more, read our next article:

What skills and knowledge do you need as a leader in the AI-era, and how will you identify, select and develop them?

A PROUD MEMBER OF AESC

PEOPLE ASSESSMENT SERVICES DNV ISO 10667-2

Reliable Leadership Advice™ MU