Nemko Digital

AI Literacy Is Now a Risk Control

Most organisations are investing in AI tools. Fewer are investing in the judgement needed to use them safely. That gap is becoming a material risk.

Across sectors, businesses are deploying copilots, chatbots, automated workflows, decision-support systems, and generative AI tools at increasing speed. Yet in many organisations, governance efforts remain focused on procurement, model performance, and compliance documentation. These are important foundations, but they are only part of the picture.

A growing share of AI risk now sits in how people understand, trust, and use these systems in practice. At Nemko Digital, we see that when employees over-rely on outputs, when managers assume automation means accuracy, or when users do not understand the limitations of AI-enabled products, organisations face operational, legal, and reputational exposure. In this environment, **AI literacy is no longer simply a learning initiative**, but it is becoming a core risk control.

AI adoption is moving faster than AI understanding

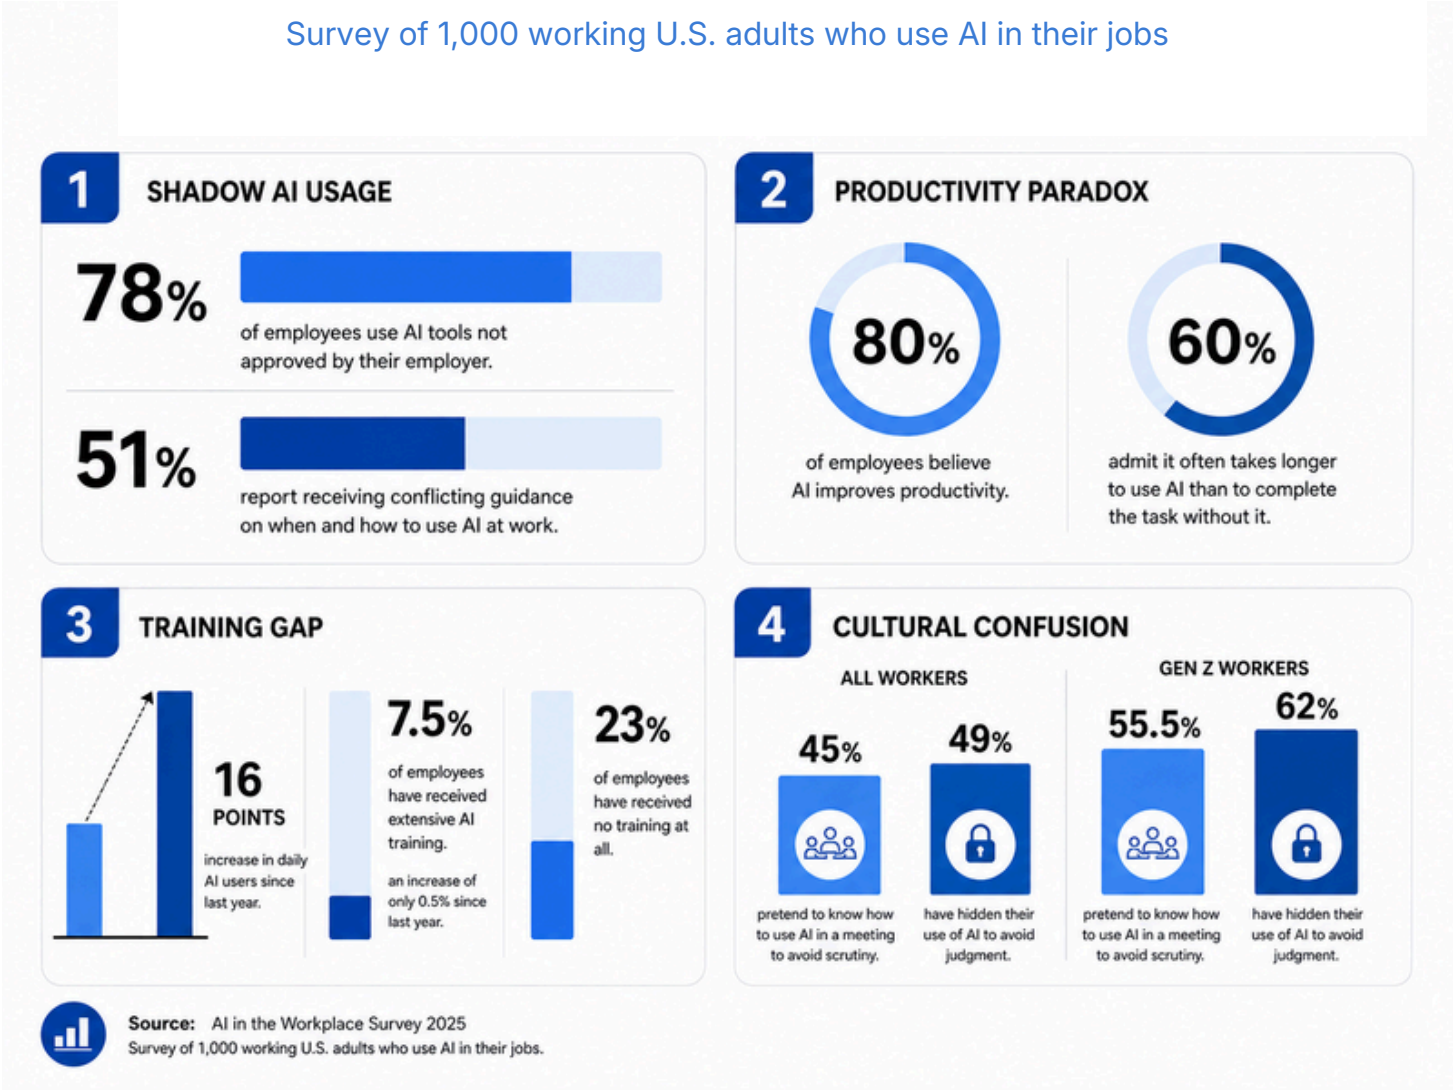
Many organisations now have access to AI tools before they have developed the internal capabilities to use them responsibly. This creates a familiar pattern: technology maturity outpaces organisational maturity.

Employees are increasingly encouraged to experiment with generative AI, often without clear guidance or oversight. In practice, this has led to the rise of "shadow AI", i.e. the use of tools outside formal governance structures.

This is already happening at scale. A 2025 SAP report found that 78% of employees use AI tools not formally approved by their employer, highlighting how adoption is often driven bottom-up rather than through controlled rollout.

AI in the workplace survey 2025

Survey of 1,000 working U.S. adults who use AI in their jobs



At the same time, accountability gaps are becoming visible when AI is embedded into workflows. In 2025, Deloitte Australia was forced to partially refund a government contract after an AI-assisted report included fabricated references and misattributed content.

Together, these examples illustrate a consistent pattern: AI adoption is often driven by speed and perceived productivity gains rather than structured governance.

Without clear guidance and accountability, organisations risk embedding these issues at scale.

This does not mean organisations should slow innovation. It means adoption should be matched with the practical judgement needed to manage risk and realise value. The challenge is not only whether an AI system works, but also whether the surrounding organisation is ready to use it well.

Why lack of AI literacy is a risk issue

AI literacy is often framed as training. In reality, it is increasingly a governance matter. Low levels of literacy can create risks such as:

- Inaccurate outputs being used in business decisions
- Sensitive data being entered into unapproved tools
- Staff relying on summaries without verification
- Human reviewers failing to challenge weak recommendations
- Poor purchasing decisions based on inflated vendor claims
- Unclear ownership when AI-assisted decisions go wrong

These are not theoretical concerns. They emerge in everyday business processes, often through small decisions made at scale. In one recent case, an AI coding agent powered by Anthropic's Claude deleted an entire production database and backups within seconds while executing a routine task, highlighting how quickly insufficient oversight of autonomous systems can translate into irreversible operational impact.

For this reason, **organisations should treat AI literacy in the same category as other enabling controls:** awareness, competence, clear responsibilities, and effective oversight.

The strategic challenge for established companies

Many AI governance programmes focus heavily on the technology itself: models, data, validation, documentation, and controls. These remain essential. But they are not sufficient.

Risk often arises in the interaction between humans and systems. For example:

- A user trusts a confident response more than a cautious one
- A manager assumes AI outputs are objective because they are data-driven
- A team adopts a tool informally outside approved process
- A reviewer signs off recommendations without meaningful scrutiny

This is sometimes described as automation bias or overtrust, but the broader issue is behavioural governance. Responsible AI requires not only responsible systems, but also their responsible use.



AI Literacy in regulation

This shift is also visible in regulation. **The EU AI Act** introduces obligations across the AI lifecycle for providers, deployers, importers, and distributors, depending on their role and the level of risk involved. It also **explicitly references the need for a sufficient level of AI literacy** among relevant persons dealing with AI systems.

That matters. It signals that regulators increasingly recognise that trust in AI is not created through technical controls alone. Competence, understanding, and informed human oversight are part of the compliance landscape.

For organisations, the practical message is clear: literacy should not be treated as optional culture-building. It should be considered part of a **defensible governance framework.**

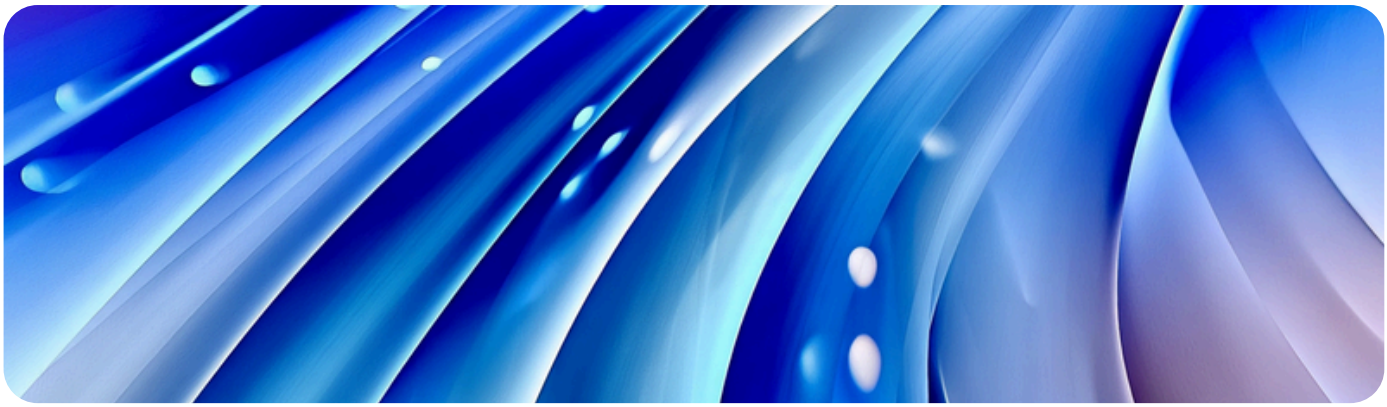
When is 'good enough'?

A common question for organisations is whether everyone now needs to become an AI expert. The answer is no.

AI literacy should be proportionate to role, responsibility, and risk. Good enough does not mean universal technical depth. It means people have enough understanding to perform their role safely and effectively. For example:

- A general employee using low-risk productivity tools should understand usage guidance, recognise data handling requirements, and understand system limitations.
- A manager using AI-supported analytics should challenge outputs, exercise escalation judgement, and take accountability for decisions.
- Teams operating higher-risk use cases should govern, validate, and oversee AI outputs with greater rigour.
- Procurement functions should assess supplier claims, evaluate controls, and review supporting evidence before adopting AI solutions.

The right benchmark is not perfection, but fitness for purpose.



Literacy and transparency: different but connected

AI literacy is often confused with transparency. They are related, but distinct.

Transparency concerns what organisations communicate about AI systems: intended use, limitations, involvement of automation, performance boundaries, or the need for human review. Literacy concerns whether people can understand and act on that information appropriately.

Transparency provides information. Literacy provides the capability to use that information well. This distinction matters especially for external users and customers. If organisations sell AI-enabled products or services, they should consider what business buyers or consumers need in order to use those solutions responsibly. That may include:

- Clear communication when AI is involved
- Understandable descriptions of capabilities and limits
- Guidance on appropriate human oversight
- Warnings against unsuitable use cases
- Clear routes for support, challenge, or escalation where needed

Transparency without literacy can become box-ticking disclosure. Literacy without transparency forces users to guess. Trustworthy AI needs both.

What is sufficient evidence?

As governance expectations mature, many organisations will face a practical question: **how do we demonstrate meaningful AI literacy?**

There is unlikely to be a single universal checklist. However, credible evidence may include:

- Role-based AI training programmes
- Completion records and refresh cycles
- Policies for approved use of AI tools
- Practical guidance for high-risk functions
- Leadership education on oversight duties
- Internal awareness campaigns
- Testing whether users actually understand expected behaviours

The real issue is whether literacy is embedded in behaviour and decision-making. A polished slide deck with no operational impact is weak evidence. Practical controls that shape conduct are stronger evidence.

What leading organisations are doing now

More mature organisations are moving beyond generic awareness sessions and embedding literacy into operating models. This often includes:

- Segmenting literacy requirements by role and risk level
- Aligning training to actual business workflows
- Combining policy with scenario-based exercises
- Integrating AI use guidance into onboarding and controls frameworks
- Educating executives on governance accountability
- Reviewing supplier transparency claims during procurement

Some organisations are also preparing for future regulatory guidance and market expectations by **treating literacy as part of broader AI management systems.**



Why it matters economically

AI literacy is not only about reducing downside risk. It is also a business enabler that can accelerate value creation from AI investments. **BCG finds** that organisations generating meaningful returns from AI are those that invest systematically in upskilling their workforce, while **limited AI literacy remains a primary factor behind the gap between ambition and realised value.**

When employees and decision-makers **understand how to use AI effectively**, organisations are better positioned to adopt new tools with confidence, integrate them into workflows more efficiently, and achieve stronger output quality. Teams are more likely to challenge weak results, apply appropriate human oversight, and use AI where it adds genuine value rather than creating unnecessary friction.

This can translate into fewer operational incidents, lower rework costs, and more informed procurement decisions when selecting external AI solutions. It can also **strengthen client confidence**, as customers increasingly expect suppliers to demonstrate responsible and competent use of AI.

From a regulatory perspective, organisations with stronger internal literacy are often **better prepared to respond to governance expectations**, evidence requests, and evolving compliance obligations.

By contrast, l**ow literacy can create hidden drag across the business**. Uncertainty slows adoption, misuse creates avoidable issues, duplicated controls increase cost, and poor outcomes undermine confidence in future AI initiatives.

Ultimately, trustworthy AI adoption depends as much on human capability as it does on technical capability.

From AI capability to AI readiness

The next phase of AI governance is about whether organisations are ready to use the systems responsibly, and not only about what these systems can do.

Many businesses continue to ask: Which tools should we deploy?

However, an equally important question is now emerging: **Do our people understand how to use these AI tools safely, critically, and accountably?**

Because in many organisations, one of the most important AI controls may not be technical at all. **It may be informed human judgement.**

About the author



Alicja Halbryt

Before joining Nemko Digital, Alicja Halbryt worked for the Dutch Ministry of Economic Affairs as an AI standardisation expert at CEN/CENELEC JTC 21. She holds an MSc in Philosophy of Technology and an MA in Human-Centred Design. She is dedicated to shaping an ethical and human-aligned AI ecosystem.

Nemko Digital