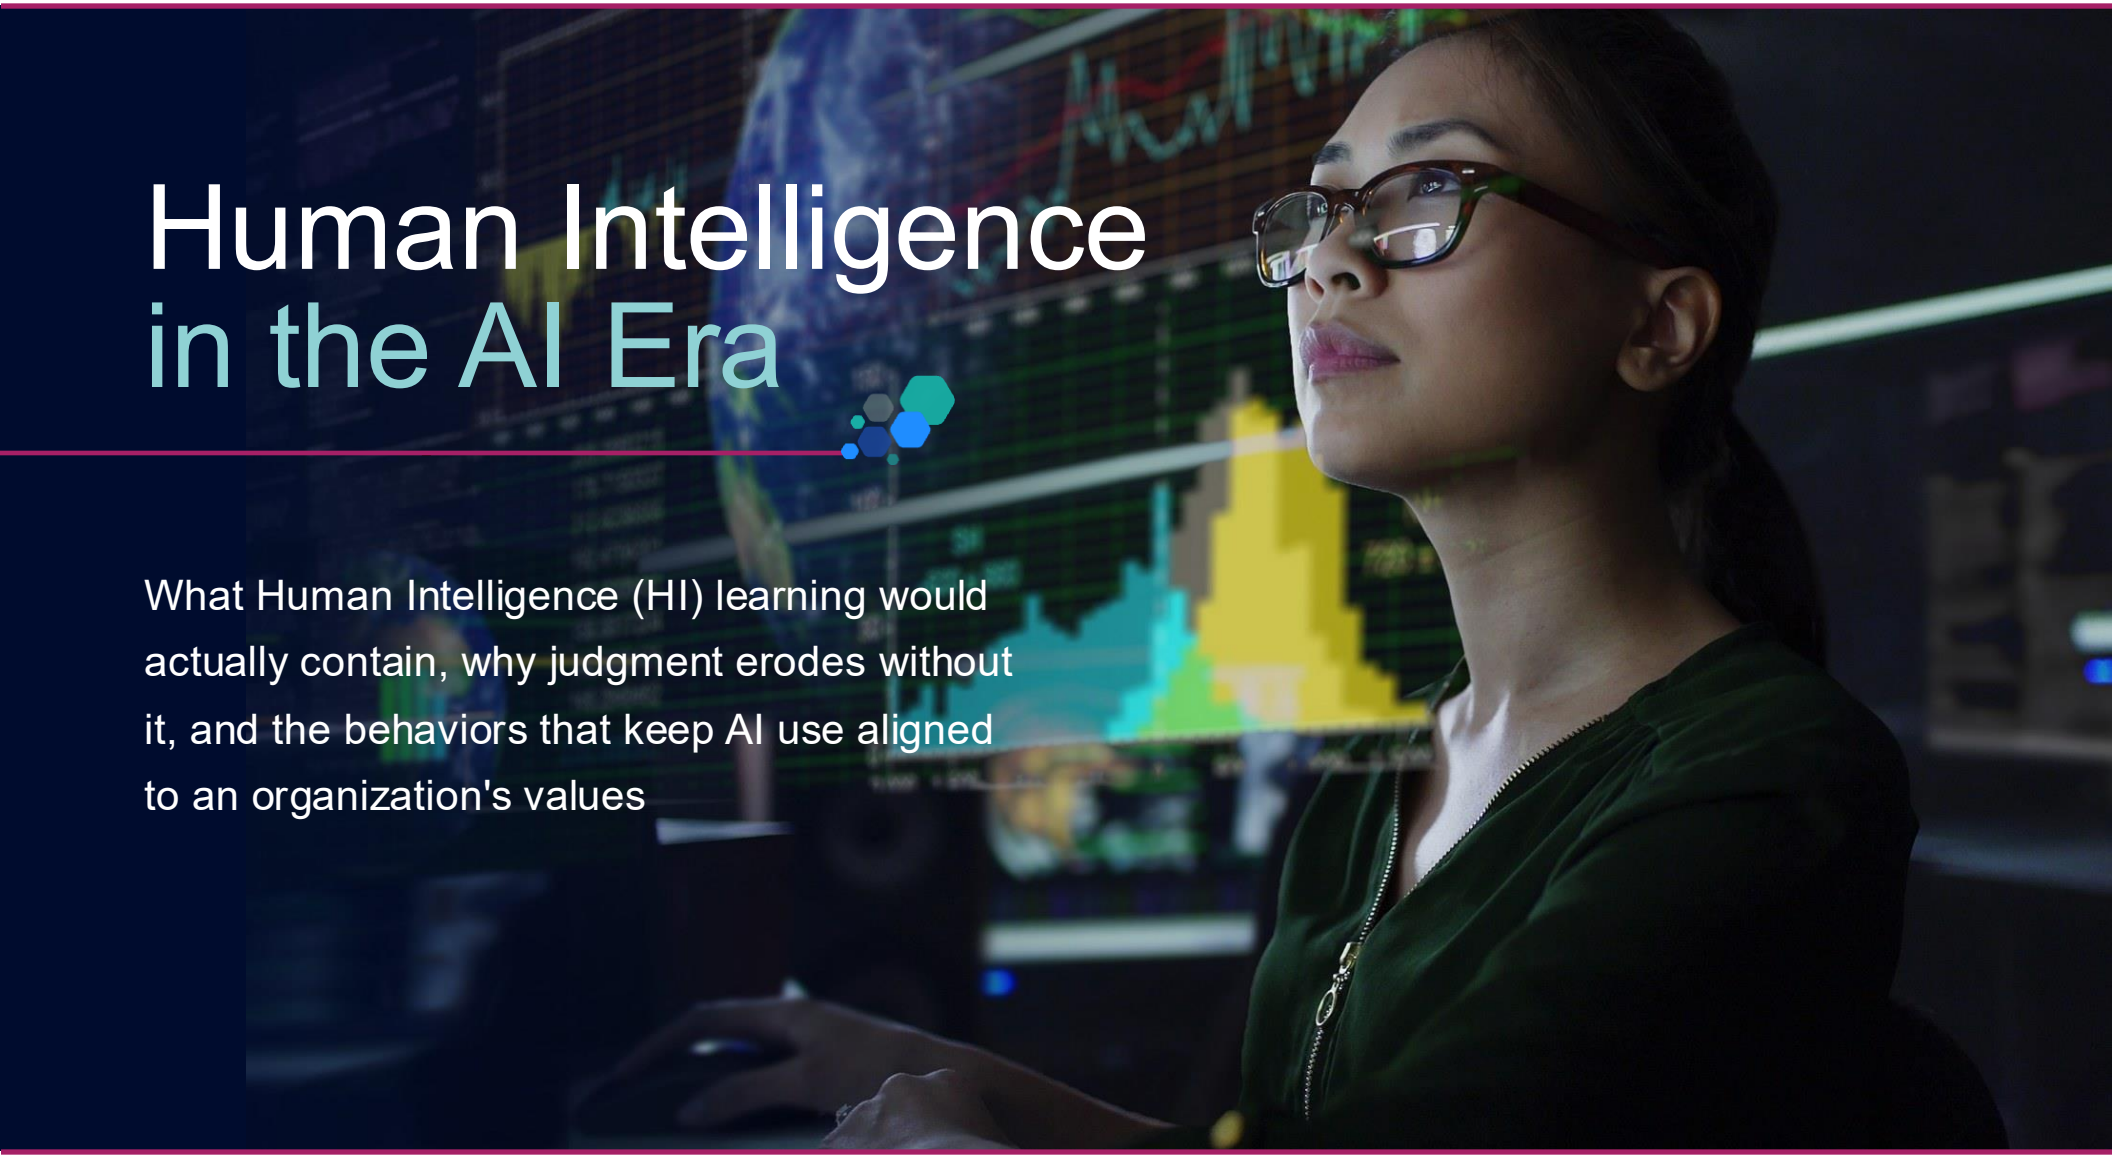
LRN Inspiring Principled Performance

Point of view · Human intelligence

October 2026

Human Intelligence in the AI Era



What Human Intelligence (HI) learning would actually contain, why judgment erodes without it, and the behaviors that keep AI use aligned to an organization's values

AI learning has a curriculum. Human intelligence learning does not. The capabilities we say matter most do not arrive as a by-product of AI adoption — under ordinary working conditions they erode, and the organizations that notice are the ones that develop them on purpose.

Executive summary

The argument in brief

Most organizations can describe their AI learning in detail: prompting, tool selection, data handling, what not to paste into a public model. Ask the same organizations what their human intelligence learning contains, and the answer is a list of qualities rather than a program.

Those qualities will not develop on their own. Working alongside capable AI does not automatically strengthen the capabilities we say we want. Under ordinary conditions it weakens them, in three distinct ways: confidence in the system lowers scrutiny, experience is not a defense, and the underlying skill decays with exposure — measurably, within months.

HI learning is a set of behaviors, not a list of virtues: verification, ethical judgment under pressure, decision-surfacing, drawing insights from data, deliberate delegation, ownership, and translation. Every one of them can be taught, practiced and observed, and none of them requires new technology.

The working day decides whether they survive. What the job does not reinforce, people stop doing. Anything that slows the work down enough to make people think will be unpopular, however well it works. That is exactly why it needs protecting — and why leaders who visibly practice these behaviors matter more than any policy about them.



AI learning has a curriculum. HI does not.

Earlier this year, at a private LRN roundtable of chief human resources officers, a participant recalled a question a senior leader had put to them.

“You're doing all this AI training. Are you doing HI training?”

HI: human intelligence. The distinctly human capabilities that become more valuable as AI becomes more capable — critical thinking, ethical judgment, curiosity, empathy, communication, collaboration, and the ability to navigate ambiguity.

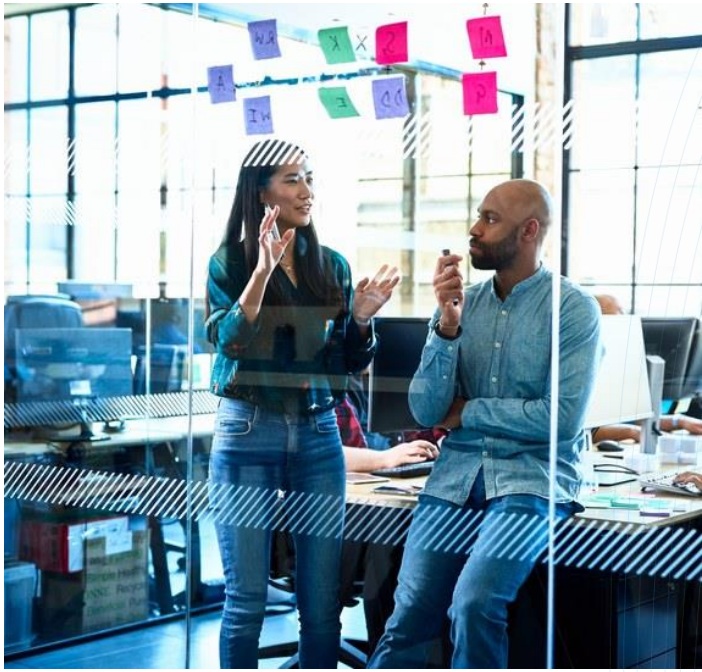
The question was easy to agree with. It has turned out to be much harder to answer.

A curriculum on one side, a list on the other

Most organizations can now describe their AI learning in some detail. Prompting. Tool selection. Data handling. What not to paste into a public model, and where the policy sits.

Ask the same organizations what their human intelligence learning contains, and the answer tends to be a list of qualities rather than a program. We are clear about what we want people to be. We are much less clear about what we want them to do differently.

That gap matters, because the qualities on the list will not develop on their own. The evidence of the last eighteen months says something stronger: under ordinary working conditions, they erode.



The demand side is not in doubt

In the World Economic Forum's Future of Jobs Report 2025, seven of the ten fastest-growing skills to 2030 are human ones — creative thinking, resilience and agility, curiosity and lifelong learning, leadership and social influence, and analytical thinking among them.

LinkedIn's Skills on the Rise 2026 tells the same story from live hiring data and adds one of interest: risk and compliance management is now one of the fastest-rising skill clusters in the market.

Workers can see it too. In Microsoft's 2026 Work Trend Index, the two skills people ranked as most important for the next few years were quality control of AI output, at 50%, and critical thinking, at 46%.

Judgment is becoming the job

A participant at the roundtable described the shift as a reversal. Professionals have historically spent something like 90% of their time researching and executing and 10% applying judgment. As AI absorbs more of the execution, that ratio could invert.

Or, as Wharton's Stefano Puntoni puts it: “The job of the human expert is not going to be doing things, but deciding what needs to be done.”

Adoption is ahead of understanding

LRN's 2026 Ethics & Compliance Program Effectiveness Report found that 39% of organizations already use AI in some part of their program, yet fewer than half can explain how it is improving outcomes.

If judgment is becoming the job, it needs to be resourced like one.



Three ways judgment quietly erodes

Working alongside capable AI does not automatically strengthen the capabilities we say we want. The research now shows three distinct ways it weakens them.

1 Confidence lowers scrutiny

Microsoft Research and Carnegie Mellon University found that the more confidence people placed in an AI system, the less critical thinking they applied to its output. Wharton researchers put a number on it: across nearly 10,000 reasoning trials, people followed a chatbot's answer roughly 80% of the time when it was wrong, and their confidence rose either way. The authors call it cognitive surrender.

2 Experience is not a defense

In a 2023 Radiology study, radiologists with more than fifteen years' experience read mammograms at 82% accuracy when an AI suggestion was correct and 45% when it was wrong. Less experienced readers fell below 20%. The people an organization relies on to catch the machine's errors are pulled toward them.

3 The skill itself decays

The Lancet Gastroenterology & Hepatology reported on nineteen experienced endoscopists across four centers. After a few months working with AI detection tools, their unassisted detection rate fell from 28.4% to 22.4%. Not in the moment of over-trusting a suggestion — in the underlying capability, once the tool was taken away.

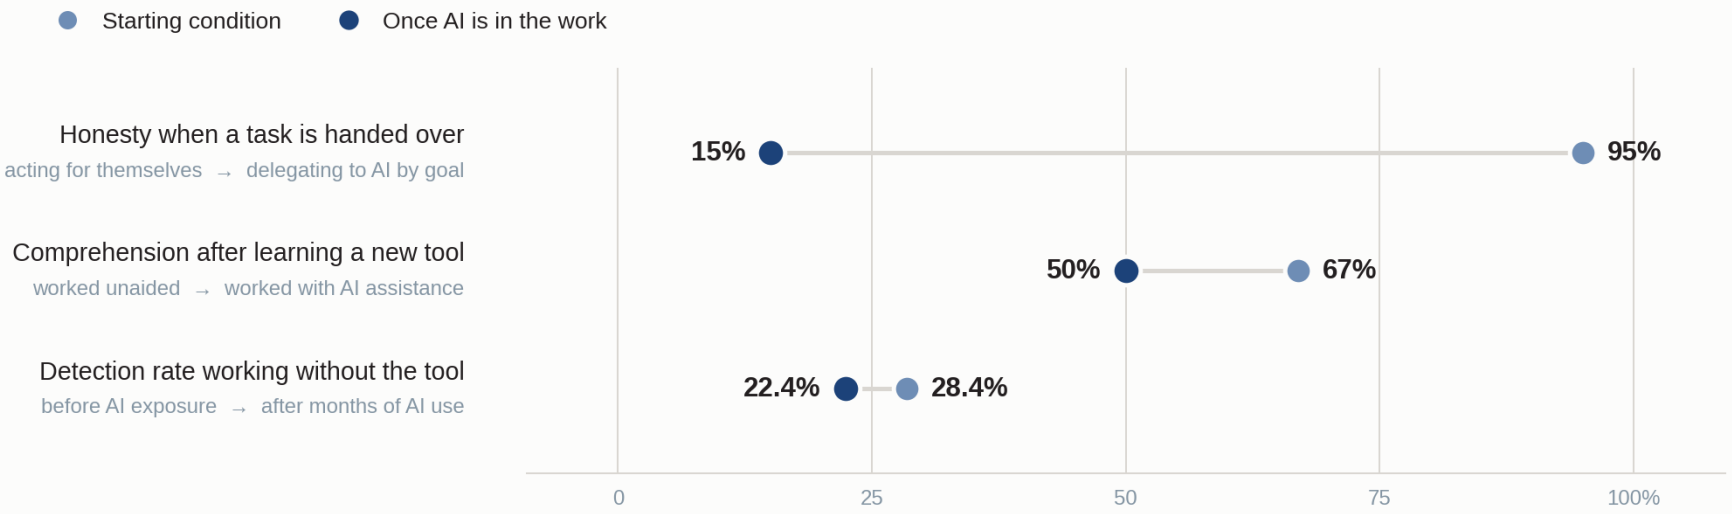
The same pattern is showing up in knowledge work. Anthropic's randomized trial of developers learning a new tool found that those working with AI assistance scored 50% on a comprehension test against 67% for those working unaided — nearly two letter grades — for a time saving of about two minutes.

Boston Consulting Group reported in June 2026 that half of the C-suite leaders it surveyed already observe de-skilling, that 90% name over-reliance on AI output without critical testing as a primary concern, and that only one in ten has an organization-wide plan to address it.

The evidence, in numbers

Three findings measure the same thing in different settings and illuminate what happens to a human capability once AI enters the work. They are different professions and different tasks, and each comparison is self-contained but the direction is the same in all three, and none of the drops is small.

Two further findings are not shown here because they are not measured on this scale. David Krakauer of the Santa Fe Institute draws the underlying distinction between cognitive tools that leave you stronger when they are taken away, and tools that leave you helpless without them. Which one AI turns out to be is not a property of the technology. It is a property of how we choose to work with it.



Honesty: Nature, 2025. Comprehension: Anthropic randomised trial, 2026. Detection rate: Lancet Gastroenterology & Hepatology, 2025.

The Anthropic result is the one to read twice: the developers who kept their understanding were not the ones who used AI less. They were the ones who asked it to explain, questioned what it generated, and used it to build comprehension rather than to skip it. How people work with the tool mattered more than whether they used it — which is precisely what makes this a learning problem rather than a tooling one.

More AI learning will not fix this

The instinct is to answer skill erosion with more AI skills. The evidence suggests that is the wrong lever.

Literacy is not judgment

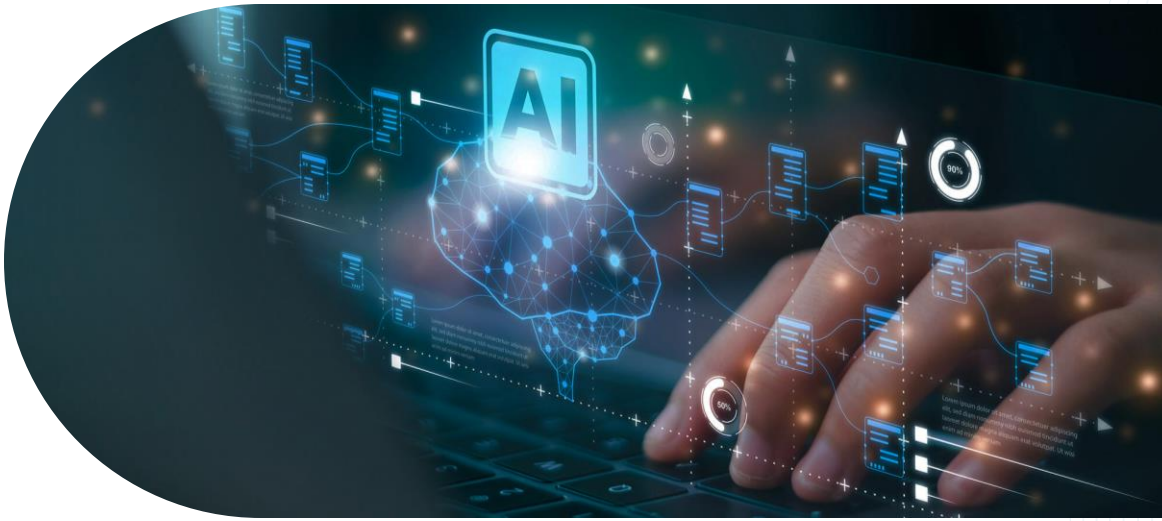
A 2025 study in the Journal of Marketing found that lower AI literacy predicts greater receptivity to AI: the people who understand the technology least find it most magical, and trust it most readily. Literacy about the tool tempers that awe. It does not by itself supply judgment about the output.

Sophistication is rare, and it is behavioral

KPMG and the University of Texas at Austin analyzed 1.4 million real workplace AI interactions across some 2,500 employees over eight months. While 90% used AI regularly, only around 5% used it with real sophistication: requiring explanations of its reasoning, setting boundaries and structure, iterating on outputs, and treating the system as a reasoning partner rather than an answer machine.

Those are not prompting techniques. They are habits of mind, and they are observable — which is the whole reason they can be developed.

AI learning teaches people how to operate the tool. HI learning is what determines what they do with the output.



What HI learning should contain

Not a list of virtues. A set of behaviors that can be taught, practiced and observed — none of which requires new technology.



Verification

Checking output against its sources, and being able to say why an answer is wrong rather than only that it feels wrong. Teachable: lateral-reading instruction lifted the share of students who verified a claim from 12% to 52%.



Ethical judgment under pressure

Recognizing when an instruction to a person, or to a model, has been framed as a goal that invites a shortcut, and escalating rather than proceeding. Goal-based delegation is where honesty collapses, from around 95% to 15%.



Decision-surfacing

Asking what the model decided on your behalf including the assumptions, the exclusions, the framing and treating those decisions as things to be settled rather than discovered later.



Drawing insights from data

Reading the data behind an answer — where it came from, what it leaves out and what it can support — and forming a view before the model forms one for you. Employers rank analytical thinking as the top core skill in the WEF's 2025 survey.



Deliberate delegation

Pausing to decide whether a task belongs to a person or a tool, and which parts of it. 53% of Microsoft's frontier professionals do exactly that, and 43% intentionally complete some work without AI to keep their skills sharp.



Ownership

"Use AI as a tool, not a substitute for judgement — check the results, because the outcome remains yours." Betty Yang, Olam Food Ingredients. The alternative, in BCG's words: "the AI suggested it" becomes accountability avoidance.



Translation

Being able to explain to a team how AI is being used, what its limits are, and how someone can challenge an outcome.

What it looks like on a Tuesday

An illustrative case, in territory ethics and compliance teams already own.

A third-party due diligence analyst uses an AI assistant to screen a new supplier. The tool returns a clean, well-structured summary: no adverse media, no sanctions matches, low risk. It reads well. It arrives in seconds.

In reviewer mode, the analyst checks the summary is formatted correctly, sees nothing alarming, and approves. Nothing about that decision looks like a failure: no error to catch, no alert, and no moment at which the analyst is asked to think.

The same task, with the behaviors

The analyst asks which sources were searched and which were not and discover the tool did not cover local-language press in the supplier's home market. They ask what the model treated as a match threshold and learn that a near-match on a director's name was discounted.

They notice the request was framed as “confirm this supplier is low risk”, a goal that invites agreement and reframe it as “list the reasons it might not be.” They run one check by hand. And when the supplier is approved, they record that they approved it, not that the tool did.



None of that required a new system, a new policy, or a procurement cycle. It took about fifteen minutes, and a set of habits somebody had taught them. Verification, decision-surfacing, drawing insights from data, ethical judgment under pressure, deliberate delegation and ownership all appear in that one screening — which is what it means to say HI is a curriculum rather than a list of qualities.

The workflow teaches more than the course

One-off learning does not hold, particularly for something moving this quickly. As LRN's Chief People & Culture Officer, Patsy Doerr, has put it: “Learning has to be ongoing and embedded into how work actually happens, and if we are not reinforcing and recognizing learning behaviors, they will not stick. Reward curiosity, adaptability and growth, not just outcomes."

Friction is the part that works, and the part nobody likes

For human intelligence there is a sharper version of that point. The working day either creates the conditions for judgment or it does not. Research on cognitive forcing functions including small, deliberate pauses that make people think before accepting an AI recommendation, finds that they reduce over-reliance, and that people rate them least favorably precisely when they work best.

Nobody enjoys friction. Which means the behaviors will not survive on individual willpower. They have to be designed into the work, and leaders have to protect them when they slow things down. The next page is about what that asks of them.

Regulators are saying the same thing

The US Department of Justice asks how quickly a company can detect and correct decisions made by AI that are inconsistent with its values. The EU AI Act goes further: those assigned to oversee high-risk systems must “remain aware of the possible tendency of automatically relying or over-relying on the output” — automation bias, named in law.

Speed of detection is not a policy. It is a property of the workflow.



Leaders set the pattern

Behaviors that slow work down survive only where leaders visibly practice them, ask for them, reward them and stop asking for the fast answer that quietly punishes deliberation.

Role-modelling, not mandating

This is a leadership skill, not a management one. Managers can require that an AI output is checked. Only leaders can show what checking looks like: asking a model to explain its reasoning in a meeting, saying which parts of a decision they kept for themselves, changing their mind when someone challenges the tool, and owning the outcome either way. People copy what leaders do with AI far more closely than what the policy says about it.

That is how ethical culture has always been built: what leaders do carries more weight than what programs prescribe. AI does not change that. It raises the stakes, because the shortcuts are now faster and better presented than ever.

What we are doing at LRN

We are building both sides of this inside our own company, and we have not separated them. Our AI literacy and fluency program gives all of our people a shared baseline in the tools — assessed rather than assumed — with a network of AI ambassadors and regular innovation workshops that keep it moving.

Alongside it, our career and leadership development curriculum is where the human side lives: judgment, feedback, coaching, decision-making and the leadership behaviors described on this page. The intent is that a leader learning to run a team is learning, in the same place, how to lead a team that uses AI wisely.

AI fluency tells our people what the tools can do. Leadership development is where they decide what they should do.



How to know if it is working

If HI learning is a set of behaviors, it can be measured the way behaviors are measured: by whether they show up.

The question to ask	What the answer tells you	Where to look
How often is an AI recommendation challenged, changed or rejected before it is acted on?	Whether review is real or nominal. A rate near zero is not a sign of a good model; it is the absence of a control.	Workflow logs; approval and exception records
When a bad output is caught, who caught it, and how long did it take?	The organization's actual speed of detection — the thing the DOJ asks about, measured rather than asserted.	Incident and near-miss reviews
What proportion of a team's work is still done unaided, and is that proportion chosen or accidental?	Whether the skill review depends on is being maintained on purpose, or quietly eroding.	Team norms; sampling of completed work
Do people describe what they decided, or what the tool produced?	Whether ownership has transferred to the system in the language people use about their own work.	Handover notes; QA conversations; appraisals

None of these numbers are hard to collect. Very few organizations collect them, which is the practical difference between having a program and having a list. They are also the numbers that make the case internally: a challenge rate, a detection time and an unaided-work proportion are the kind of evidence a board and a regulator both recognize.

Conclusion: the question, answered

HI learning is not the softer companion to AI learning. It is the part that determines whether AI use stays inside an organization's values, and whether someone notices when it does not.

The capabilities are rising in value, and that much is not disputed. What the last eighteen months of rapid AI adoption have added is the harder half of the finding: they do not arrive on their own, and the conditions most organizations are building — more output, faster, reviewed by the same people — are the conditions under which they erode.

None of the remedies is exotic. Verification, ethical judgment under pressure, decision-surfacing, drawing insights from data, deliberate delegation, ownership and translation are ordinary professional behaviors. They can be described, practiced, observed and rewarded. What they cannot do is survive a working day designed as though they were automatic, or leaders who do not practice them themselves.

That is the answer to the question asked at the roundtable. HI learning is not a list of qualities. It is a curriculum of behaviors, a workflow that makes room for them, leaders who model them, and a set of measures that show whether they are happening.

The organizations that get this right will not be the ones whose people are most comfortable with AI.

They will be the ones whose people have kept the ability to disagree with it.



Sources and notes

Origin and LRN material

LRN, “Beyond AI: Why human intelligence is becoming leadership's greatest competitive advantage” (30 July 2026) — the HI training question, the capability list and the 90/10 reversal. Patsy Doerr in Governance Intelligence, “Learning has to be ongoing” (9 April 2026). LRN, 2026 Ethics & Compliance Program Effectiveness Report. LRN, “AI accountability in action” (31 August 2026) — Betty Yang quote. LRN white paper, Knowing Is Not Doing (June 2026).

Demand for human capability

World Economic Forum, Future of Jobs Report 2025, Skills Outlook. LinkedIn, Skills on the Rise 2026 (February 2026). Microsoft, Work Trend Index 2026. Wharton Human-AI Research, Accountable Acceleration (2025) — Stefano Puntoni.

Erosion of judgment and skill

Lee et al., “The Impact of Generative AI on Critical Thinking,” Microsoft Research and Carnegie Mellon University, CHI 2025. Shaw and Nave, “Thinking — Fast, Slow, and Artificial,” Wharton working paper (January 2026). Dratsch et al., “Automation Bias in Mammography,” Radiology (May 2023). Budzyń et al., “Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy,” The Lancet Gastroenterology & Hepatology (online August 2025). Anthropic, “How AI assistance impacts the formation of coding skills” (January 2026). Boston Consulting Group, “When Everyone Uses AI, Companies Risk Losing Critical Skills” (June 2026). David Krakauer, Santa Fe Institute, on complementary and competitive cognitive artifacts (September 2025).

Literacy, sophistication and practice

Tully, Longoni and Appel, “Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity,” Journal of Marketing (2025). Kowaleski, Hallman and Schmidt with KPMG, “What the Best AI Users Do Differently,” Harvard Business Review (March 2026). Brodsky et al. on lateral-reading instruction, AERA Open (2021). Buçinca, Malaya and Gajos, “To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI,” PACM HCI (2021). Köbis et al., “Delegation to artificial intelligence can increase dishonest behaviour,” Nature (September 2025).

Regulatory

US Department of Justice, Evaluation of Corporate Compliance Programs (September 2024). Regulation (EU) 2024/1689 (EU AI Act), Article 14(4)(b) on automation bias.

Note on the figure

The three findings on page 6 are measured on a common percentage scale but in different professions and tasks; each comparison is self-contained and the rows are not intended to be compared with one another. Findings reported as a change in points against a baseline, or as a correlation, are described in the text rather than charted.



About LRN. LRN helps organizations build the culture, capabilities and systems that make ethical behavior scalable — through measurement, governance, enablement and an AI-era ethics and compliance curriculum. Its work helps leaders translate principles into practical controls for human and AI-enabled decision-making.

About this paper. This point of view synthesizes peer-reviewed research, randomized trials, industry survey data and regulatory guidance, together with LRN's own program research. Findings from controlled studies are presented as evidence that an effect can occur, not as an estimate of its frequency in any one organization. Figures are current to September 2026.

LRN Inspiring Principled Performance

2026 LRN. All rights reserved. All brand, product, service names, and logos are trademarks and/or registered trademarks of their respective manufacturers and companies.