ESCP BUSINESS SCHOOL

IT ALL STARTS HERE

BERLIN I LONDON I MADRID I PARIS I TURIN I WARSAW

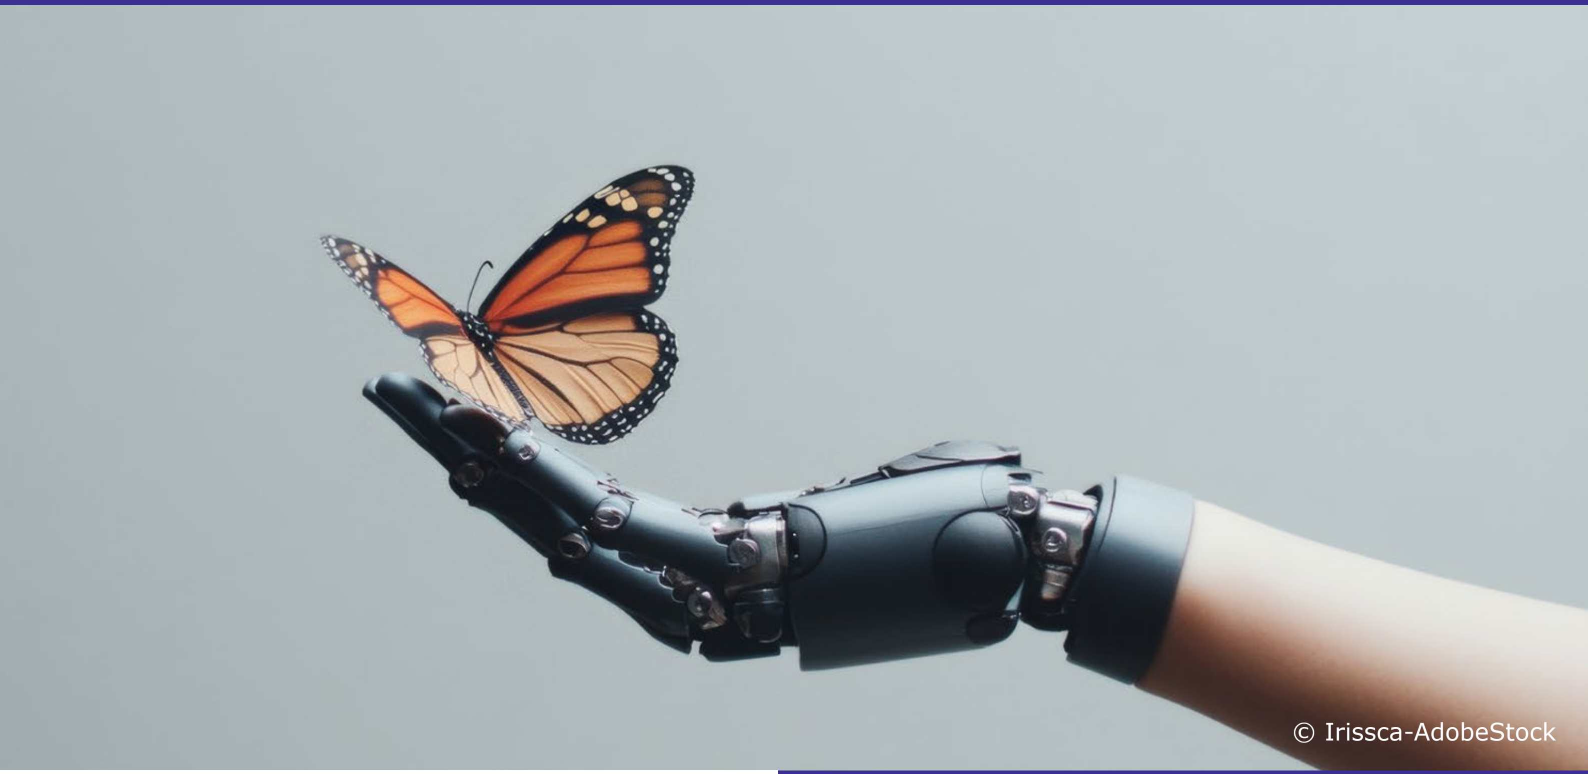

© Irissca-AdobeStock

LIGHTS - Human/Well being

What counts as intelligence? Rethinking AI and organizational knowing

ESCP Impact Paper No.2025-28-EN

Hsin-Hsuan Meg LEE
ESCP Business School

ESCP RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ESCP Impact Paper No.2025-28-EN

What counts as intelligence?
Rethinking AI and organizational knowing

Hsin-Hsuan Meg Lee*

ESCP Business School

Abstract

As artificial intelligence becomes more deeply embedded in business decision-making, organisations are redefining what counts as knowledge as they adopt new tools. This paper examines how AI systems structure the conditions of insight by prioritising what can be legibly quantified and predicted. Drawing on examples from language learning, writing support, taste modeling, and hiring, the paper shows how AI systems tend to privilege machine-compatible forms of intelligence while excluding context-based, intuitive, or relational ways of knowing. These exclusions are not simply technical limitations, but reflections of deeper epistemic assumptions. The paper argues that managers need to be more deliberate about the forms of intelligence they encode into systems, especially when building or selecting AI tools. It concludes with a set of practical questions to help managers assess what their systems are designed to recognise and what may be quietly left out.

Keywords: AI, organisational knowledge, epistemic exclusion, managerial judgement, intelligence design

*Associate Professor, ESCP Business School

What counts as intelligence? Rethinking AI and organizational knowing

Introduction: what counts as smart?

In a recent marketing class, I posed a question that sparked a lively discussion. If you were managing a pet food company preparing to launch a new product for dogs, would you consider hiring an animal communicator to understand your customers better? The idea came from the rise of viral pet psychics on TikTok (Roy, 2025). Most students laughed and quickly dismissed the idea as pseudoscience. A few said, "Why not?" and thought it could offer creative insights. Some others argued that pet owners, who had spent enough time with their dogs, could answer any relevant questions themselves.

I then asked them to reflect on whether their own observations and intuitions as pet owners felt more scientific, and why. After that, I introduced a second option. What if we used an AI system trained to analyse animal behaviours using neural signals or wearables to decode communication (e.g., Abzaliev et al., 2024)? That changed the tone in the room. Most students said "maybe." One eventually concluded that it did not matter either way, because humans are the ones buying the product. The animals' "voices," even if accessible, would not be considered valid or relevant data.

This discussion raised a central question for me. Who decides what counts as a valid source of intelligence, and in what form? As AI becomes more embedded in business decision-making, companies are not just collecting new types of data. They are also redefining what it means to know something. These systems reward legibility, pattern recognition, and machine compatibility. In doing so, they often push aside intuition, relational insight, and context-specific intelligence.

This paper explores how organisations increasingly rely on machine-readable, data-driven insights and downplay other ways of knowing. I draw on real-world examples and interdisciplinary theory to argue that AI systems do not simply process data. They reflect and reinforce assumptions about what intelligence is. Much of the current debate about AI ethics focuses on fairness, bias, or explainability. This paper takes a step back to ask a more basic question: What kinds of knowledge do we consider valid, and which ones are being ignored?

How AI systems define intelligence and what they leave out

AI systems reflect design assumptions about what matters. Marketing research has long favoured rational, measurable forms of knowledge over embodied or intuitive ones (Kravets & Varman, 2022). AI extends this preference by prioritising what is trackable and reproducible to optimise it. Ways of knowing that fall outside what is easily legible are often treated as noise (Fricker, 2007), and this pattern is now embedded in how machines process information.

For example, Duolingo's AI rewards fast recall, standardised pronunciation, and multiple-choice accuracy[1]. These metrics define language proficiency narrowly. The system cannot assess code-switching, conversational nuance, or improvisation. Learners may perform well in the app while struggling in real-world settings. By teaching what it can track, Duolingo

[1] https://blog.duolingo.com/large-language-model-duolingo-lessons/

reshapes learners' sense of what it means to know a language. This reflects a broader trend in AI: what cannot be measured is rarely treated as intelligence.

Similarly, AI writing tools like Grammarly or ChatGPT reward clear, coherent, and grammatically correct outputs. These systems shape norms around "good writing," favouring fluency, neutrality, and standard usage. Writers may begin shaping their expression to match the system's preference, abandoning irony, cultural reference, or creative complexity. Over time, intelligence becomes aligned with what the system can polish, not what the writer can explore.

Stitch Fix builds customer profiles based on behavioural data such as clicks, purchases, and returns.2 These patterns are used to infer style preferences and suggest clothing that aligns with prior choices. But taste is not only behavioural. It is also intuitive, expressive, and shaped by culture and our own contextual identity. These qualities are more complex to model and are often excluded from algorithmic logic. Aesthetic intelligence is reduced to pattern detection, flattening personal expression into algorithmic predictability.

AI-assisted hiring tools present a further case. These systems often evaluate candidates based on forms of expression such as tone, speech rate, and fluency. These features stand in for emotional or interpersonal intelligence but rely on a narrow model. True emotional intelligence involves attunement, contextual response, and the ability to handle ambiguity, qualities that are often illegible to systems trained on surface features. As Hunkenschroer and Luetge (2022) note, such tools may not just filter candidates but also redefine what counts as a communicative or emotionally competent person.

Across these examples, AI defines intelligence in terms of what can be modelled and measured. It tends to overlook what is uncertain, felt, or embedded in context, the very qualities that often shape learning and judgment.

Framing the problem: intelligence as a design decision

Current AI systems are designed around assumptions that treat intelligence as symbolic, extractable, and abstract. Yet human and more-than-human forms of intelligence often rely on embodied, sensory, and affective ways of knowing that emerge through physical interaction and contextual awareness. These qualities are difficult to formalise within training data or model architectures. As Paolo et al. (2024) argue, engaging seriously with embodied intelligence requires moving beyond current design assumptions.

Once embedded in decision-making, AI systems begin to shape what counts as knowledge. Most commercial models are trained on data that is digitised, standardised, and drawn predominantly from English-language, Western sources. This limits what systems can recognise and respond to. As a result, ecological, spiritual, communal, and embodied forms of knowledge are routinely excluded (de Sousa Santos, 2024).

Even large language models, often praised for their generative flexibility, reproduce these limitations. LLMs may appear fluent across languages, but their responses often rely on translation-like shortcuts that bypass contextual or cultural nuance. This leads to a flattening of linguistic nuance and a broader loss of epistemic diversity (Zhang et al., 2023).

As organisations increasingly rely on AI to process information and define relevance, these systems begin to reconfigure how intelligence is understood. They influence which perspectives are treated as credible, what kinds of insight are valued, and how learning is

[2] https://algorithms-tour.stitchfix.com/

structured. The issue is not only about fairness or inclusion but about the narrowing of knowledge around what is easy to model.

Why this matters for managers: intelligence is strategic

For managers, the key challenge is understanding how AI shapes what their organisation treats as intelligence. Most AI tools are built for standardisation and efficiency. They perform well with structured inputs and repeatable patterns. However, many valuable insights in business are intuitive, context-dependent, or relational. As these systems become part of daily decision-making, they begin to influence which forms of intelligence are made visible and which are overlooked. This influence extends beyond individual decisions. AI tools help define how teams learn, how problems are framed, and which forms of contribution are considered credible. From hiring tools that prioritise fluency to dashboards that favour what is already legible, these systems shape internal norms around thinking clearly, leading effectively, or making a good decision.

Managers are often encouraged to balance quantitative analysis with personal judgment or intuition. But as AI systems automate more of the process, this balance becomes harder to maintain. Studies have shown that integrating AI with human insight requires conscious effort, particularly in settings where formal data is limited or ambiguous (Vincent, 2021). As the line between augmentation and automation becomes increasingly blurred, managers must actively decide which forms of judgment to retain and support (Raisch & Krakowski, 2021). Some systems now claim to incorporate more complex inputs, such as emotion, behaviour, or planetary systems. But they still reflect inherited assumptions about what counts as data and what can be known. Most forms of field-based or experiential knowledge remain outside the frame.

Managers involved in selecting or deploying AI tools should ask what the system is designed to recognize and what it overlooks. Systems reflect the data and design assumptions on which they are built. If these assumptions go unexamined, they can begin to shape culture, training, and organisational learning in subtle but lasting ways. AI can be a powerful tool for sensemaking. But its value depends on how well it matches the complexity of the world it is meant to support. Managers who are aware of what their systems include and what they leave out will be better equipped to make decisions that are efficient, thoughtful, and responsive to the full range of knowledge their organisation holds.

Conclusion: reclaiming intelligence as a design choice

The question of what counts as smart is no longer abstract. It is built into organisations' tools to decide, predict, and plan. As AI becomes more central to business operations, it reshapes how insight is defined and whose knowledge is trusted. The key issue is not whether AI systems are ethical by design, but whether organisations are deliberate about what forms of intelligence they validate. What is not recognised by the system often does not enter the decision. Over time, this can narrow what problems get noticed, whose input is valued, and how learning is structured.

Before we talk about humanising AI, we need to consider how AI has already narrowed what we mean by intelligence. Managers do not need to solve this completely; they do need to understand what their systems are doing. Here are four practical questions managers can use to guide AI implementation:

1. What kinds of intelligence is this system designed to recognise?

Review what is being measured, modelled, and scored. Ask what is missing and whether that matters to your context. For example, Hugging Face's SHADES benchmark evaluates how models represent racial identity across languages and contexts, highlighting how AI systems often encode normative assumptions into classification processes (Mitchell, cited in Heikkilä, 2025).

2. What forms of insight are hard to formalise but still relevant?

Consider whether lived experience, frontline observations, or cultural nuance should play a larger role. Intuitive judgment, cultural nuance, and context-bound expertise may not appear in dashboards but still shape meaningful outcomes. Useful methods include story-based evaluation, feedback interviews with frontline staff, and collaborative review workshops. The United Nations Development Programme (UNDP)[3], for instance, uses collective intelligence approaches that combine community insight with data modelling to address complex, multi-stakeholder challenges.

3. How does this system shape learning inside the organisation?

Look not just at what decisions are made, but how the system influences team thinking over time. If dashboards or AI outputs are redefining what counts as "good" judgment or "clear" writing, build in regular checkpoints to assess what's being prioritised or quietly devalued. Internal decision reviews or side-by-side comparisons of human and system recommendations can help teams stay aware of how their learning is being shaped.

4. Who decides what gets included as valid data?

Review who defines the categories, selects the training inputs, and sets the success metrics. If this happens in a closed loop, open it up. Mozilla's Data Futures Lab supports experimental data stewardship models that foreground community-defined values, such as indigenous language preservation and ethical metadata practices, showing how organisations can rethink what counts as data and who gets to decide[4].

AI systems reflect choices. The most effective organisations will be the ones that make those choices deliberately.

References

Artem Abzaliev, Humberto Pérez Espinosa, Rada Mihalcea. Towards Dog Bark Decoding: Leveraging Human Speech Processing for Automated Bark Classification. Submited to arXiv, 2024 DOI: 10.48550/arXiv.2404.18739

de Sousa Santos, B. (2024). AI and the epistemologies of the South. *Journal of World-Systems Research*, 30(2), 635–645.

Fricker, M. (2007). "Epistemic injustice: Power and the ethics of knowing", Oxford University Press.

[3] https://feature.undp.org/new-architecture-for-public-institutions/
[4] https://foundation.mozilla.org/en/data-futures-lab/

Heikkilä, M. (2025, June 19). *Margaret Mitchell: Artificial general intelligence is* 'just *vibes and snake oil'*. *Financial Times*.
https://www.ft.com/content/7089bff2-25fc-4a25-98bf-8828ab24f48e

Hunkenschroer, A. L., & Luetge, C. (2022). Ethics of AI-enabled recruiting and selection. *Journal of Business Ethics*, 178(2), 535–553.

Kravets, O., & Varman, R. (2022). Hierarchies of knowledge in marketing theory. *Marketing Theory*, 22(1), 9–30.

Paolo, G., Gonzalez-Billandon, J., & Kégl, B. (2024). A call for embodied AI. arXiv preprint arXiv:2402.03824.

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210.

Roy, J. (2025, May 6). You Talk to Your Pets. They Help Your Pets Answer. *The New York Times*.
https://www.nytimes.com/2025/05/06/style/pet-animal-communicators-psychic-tiktok.html

Vincent, V. U. (2021). Integrating intuition and artificial intelligence in organisational decision-making. *Business Horizons*, 64(4), 425–438.

Zhang, Y., Fan, Y., & Zhao, J. (2023). When do LLMs say the F-word? Measuring hallucination in multilingual generative tasks. arXiv preprint arXiv:2305.16339.
https://arxiv.org/abs/2305.16339