

COGNITIVE DEBT™

*The Hidden Cost of Thinking Too Fast
in the Age of AI*

The growing gap between the speed of AI adoption and the readiness of human judgment to govern it — and a practical framework for closing it.

02 Cognitive Debt, and the Four Disciplines

03 The Evidence Base — Core Research

04 The Evidence Base — Data & Governance Frameworks

05 About Trust AI™

Cognitive Debt, and the Four Disciplines

In 2025, MIT Media Lab researchers ran an EEG study measuring brain activity in people writing with and without AI assistance, and found measurably reduced neural engagement¹. They called what they measured "cognitive debt." This brief extends that finding into a governance lens: as AI becomes more capable, the value of human judgment does not shrink — it increases. The organisations that win the next decade won't be the ones with the most powerful AI. They'll be the ones that preserve the strongest human judgment inside it.

Four disciplines, drawn from professions that already operate under exactly this pressure, hold the practical counter-force.

01 Clear — *the chess player's filter*

Expertise is a better filter, not a bigger engine: experts reach strong decisions faster through pattern recognition, not by processing more information⁵.

In practice: before approving any AI-generated output, write one sentence — "The actual decision I am making is ____" — and list what you haven't validated.

02 Decide — *the surgeon's commitment*

Naturalistic decision-making under time pressure and incomplete information is a well-studied skill in its own right⁶ — closing a decision loop is a discipline, not a compromise.

In practice: identify the open decisions you are carrying. Close one this week. Record the one assumption that could prove it wrong.

03 Detect — *catching the wrong note early*

Automation bias — placing too much trust in automated recommendations — is one of the best-documented failure patterns in human-factors research². Detection has to be built, not assumed.

In practice: build one review checkpoint where a second person challenges the reasoning behind an AI-assisted decision, not just the output.

04 Install — *the glazier's frame*

Deskilling under partial automation is a structural risk, not a character flaw⁴ — good judgment has to be built into upstream habits, not summoned in the moment⁷.

In practice: create one organisational habit that makes thoughtful review unavoidable. Treat AI literacy as an ongoing discipline, not a one-off module — an expectation now reflected in the EU AI Act¹³.

Sources Behind This Talk

Cognitive Debt™ is a governance lens, built on existing human-factors research and on the frameworks organisations already work within — not a claim to a new scientific finding of its own.

CORE RESEARCH

- 1 Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X. H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task." MIT Media Lab, arXiv: 2506.08872.
EEG study finding reduced neural connectivity in LLM-assisted writing — the origin of the term "cognitive debt" in empirical research.
- 2 Parasuraman, R., & Manzey, D. H. (2010). "Complacency and Bias in Human Use of Automation: An Attentional Integration." *Human Factors*, 52(3), 381–410.
The foundational review establishing automation bias and complacency as attentional phenomena, found in both novice and expert operators.
- 3 Risko, E. F., & Gilbert, S. J. (2016). "Cognitive Offloading." *Trends in Cognitive Sciences*, 20(9), 676–688.
Defines cognitive offloading as the use of physical or digital action to reduce internal cognitive demand, and reviews its costs and benefits.
- 4 Bainbridge, L. (1983). "Ironies of Automation." *Automatica*, 19(6), 775–779.
The classic paper on deskilling: when a system automates everything except the hardest judgment calls, the human loses the practice needed to handle exactly those calls well.
- 5 Chase, W. G., & Simon, H. A. (1973). "Perception in Chess." *Cognitive Psychology*, 4(1), 55–81; building on de Groot, A. D. (1946/1965).
Chunking theory: expertise is explained by rapid pattern recognition built from practice, not by searching more options. (De Groot's original claim that masters search fewer moves than club players has since been contested; the pattern-recognition finding is the well-supported part.)
- 6 Klein, G., Calderwood, R., & Clinton-Cirocco, A. (1985/2010). "Rapid Decision Making on the Fire Ground." *Journal of Cognitive Engineering and Decision Making*, 4(3), 186–209; Klein, G. (1998). *Sources of Power: How People Make Decisions*.
The Recognition-Primed Decision model of naturalistic decision-making under time pressure, developed from fireground commanders and since applied across high-stakes professions including medicine and the military.
- 7 Gawande, A. (2009). *The Checklist Manifesto: How to Get Things Right*. Metropolitan Books.
The case for building good judgment into upstream structure and habits, drawn from aviation and surgical checklists.

DATA POINTS

- 8 Microsoft & LinkedIn (2024). *2024 Work Trend Index Annual Report: "AI at Work Is Here. Now Comes the Hard Part."*
75% of global knowledge workers report using generative AI at work.
- 9 Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025). *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025*. The University of Melbourne and KPMG.
48,000 respondents across 47 countries. In Australia, 57% of employees rely on AI-generated output without evaluating its accuracy, and 59% report AI-related work mistakes as a result.
- 10 Shaw, S. D., & Nave, G. (2026). *"Thinking Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender."* Working paper, The Wharton School, University of Pennsylvania (SSRN 6097646).
1,372 participants, 9,593 trials. Participants accepted correct AI answers ~93% of the time; when the AI was deliberately wrong, they still accepted its answer ~80% of the time, while reporting higher confidence regardless of accuracy.
- 11 Budzyń, K., Romańczyk, M., Kitala, D., et al. (2025). *"Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study."* *Lancet Gastroenterology & Hepatology*, 10(10), 896–903 (the ACCEPT trial).
The first real-world clinical evidence of AI-induced deskilling in medicine: adenoma detection rate at non-AI-assisted colonoscopies fell from 28.4% to 22.4% among experienced endoscopists after routine AI exposure.

GOVERNANCE & REGULATORY FRAMEWORKS

- 12 UK Financial Reporting Council (2026). *Generative and Agentic AI Guidance: Risks, Mitigations and Illustrative Examples* (30 March 2026).
"Regulatory accountability for the deployment of AI tools and the quality of audit outputs remains unchanged. ... the human auditor is always accountable."
- 13 Regulation (EU) 2024/1689 (the EU AI Act), Article 4 — AI literacy. In application since 2 February 2025.
Amended by the 2026 Digital Omnibus from an obligation to "ensure" a sufficient level of AI literacy to an obligation to "support the development of" it — an obligation of effort, not of result. (The Digital Omnibus separately postponed the Act's high-risk system obligations to December 2027 / August 2028; Article 4 itself was not affected.)
- 14 APRA Prudential Standard CPS 230 — Operational Risk Management (in force 1 July 2025).
Requires APRA-regulated entities to maintain effective internal controls and personal accountability for operational risk, including for outsourced or automated processes.
- 15 ISACA COBIT; OECD AI Principles (2019, updated 2024); NIST AI Risk Management Framework 1.0 (2023); ISO/IEC 42001:2023.
Every major AI governance framework in current use assumes a thinking, accountable human in the loop — none endorses approval without understanding.

Cognitive Debt™ does not compete with any of the frameworks above — it names the risk they all rely on a human to catch: judgment that looks present, but is no longer doing the work.

ABOUT THIS BRIEF

Trust AI™

Trust AI™ is Anam Munshey Thomas's publishing and thought-leadership platform on AI governance and human judgment. Anam is a Senior Client Executive in Managed Services at EY, with 18+ years across IBM and EY spanning financial services, energy, insurance and superannuation across Australia, Singapore, New Zealand and the Middle East.

Trust AI™ Executive Brief Series

No. 001 — Cognitive Debt™, prepared for ISACA Melbourne, 20 August 2026. Each brief in this series takes one governance question raised by AI adoption and grounds it in the practical evidence behind it.

Judgment is becoming the competitive advantage of the AI era.

Stay Human in the Age of AI.

Anam Munshey Thomas · Founder, Trust AI™

trustai.au · [linkedin.com/in/anammunshey](https://www.linkedin.com/in/anammunshey)