

Words Gooding: The Triumph of the Humanities

INTS1301 TECHNOLOGY AND SOCIETY: FROM PLATO TO NATO — WEEK 11

Brian Ballsun-Stanton

Macquarie University

2026-05-19

Who's driving this?

TODAY'S QUESTION

Knowledge has become cheap; judgement has not. The final call (ship or scrap, invest or walk away) still hinges on a felt sense for risk, opportunity, and feasibility. ([Jones, 2025](#))

- How you use the tool determines what it does. Every layer (model, harness, prompt, modality) shapes the output.
- A bad result is sometimes the model, sometimes the setup, sometimes you; the skill is telling them apart.
- The same tool gives different answers in different hands. That is the tool, not a malfunction.
- Today: where is your judgement, and who is holding it?

Locus of control.

WHO DECIDES?

- **Locus of control** ([Rotter, 1966](#)): where you believe outcomes come from. Internal means your effort, choices, skill. External means luck, others, the tool itself.
- **AI-Locus of Control:** who is shaping what the AI gives you? You, or the AI? ([Torrington et al., 2025](#))
- Four tensions where this plays out: **Agency, Expertise, Embodiment, Praxis** ([Ballsun-Stanton & Torrington, 2025a](#)).
- Depending on how we think about and use AI, we can be somewhere on either side. Do we consider it to be more proficient than us, to know more? Do we consider it software that we can learn how and when to use?

The Continuum.

AI-LOCUS OF CONTROL

[Open the AI-LOC Continuum diagram \(full size\)](#)

It's all just text. So is the lever.

AI-LOCUS OF CONTROL: AGENCY

Agency: *Does the AI have an agency, or does the human retain it? ([Ballsun-Stanton & Torrington, 2025a](#))*

- LLMs do not read your words. They read tokens, atomic units shaped by the model's training. Your text gets converted first.
- Not a search engine. A “calculator for words” ([Willison, 2023](#)). Treating it like search hides capabilities AND dangers.
- How you write each prompt of a conversation shapes what comes back. If your answer arrives in the first response, you are not thinking enough. Even for cooking recipes.

Always confident. Usually correct.

AI-LOCUS OF CONTROL: EXPERTISE

Expertise: *Better than human or an input to human judgement?* ([Ballsun-Stanton & Torrington, 2025a](#))

- The drunk tutor: LLMs sound equally confident whether right or wrong. Tone tracks fluency, not accuracy ([Kudina et al., 2025](#)).
- From our longitudinal study, a student on AI-generated legal cases: *“The quote is beautiful. I’m going to use this. And then I was like, it doesn’t exist.”* ([Torrington et al., 2025](#))
- External LOC has the output as deeply founded in other-peoples-knowledge. Internal has it as the statistical consequence of our words and knowledge. (We should use good words, then.)

Try. Notice. Adjust. Repeat.

THE METHOD

“Changing the way I spoke about it [the prompt] changed the output.” ([Torrington et al., 2025](#))

- Define your goals. What are you actually trying to do?
- Choose. Right tool? Right model? Are you in a position to decide?
- Form an expectation. What do you intend to happen?
- Prompt intentionally. Why these words? How does it know what you want it to know?
- Evaluate. Edit the prior prompt. Expect a long conversation.
- Build taste. If you cannot tell good output from bad, stop.

The jagged frontier.

AI-LOCUS OF CONTROL: PRAXIS

Praxis: Automation versus augmentation? ([Ballsun-Stanton & Torrington, 2025a](#))

“I am confident in my ability to judge when AI comes in handy.” ([Torrington et al., 2025](#))

- *“Within this jagged frontier, AI can complement human work. However, outside the frontier, AI output is inaccurate, is less useful, and can degrade human performance.”* ([Dell’Acqua et al., 2026](#))
- Sometimes the right answer is wrong tool (use a citation manager), wrong model (try a different one), wrong setup (more context). Sometimes it is don’t use AI.
- Contrast: *“What’s the point of learning if AI can just do it better? You pretend to teach us, we pretend to learn.”* v. *“Ooo, now I can do this cool new thing!”* Knowing when and where to use what tool is important.
- Outsourcing the thinking accumulates cognitive debt ([Storey, 2026](#)).

No kitten in the mirror.

AI-LOCUS OF CONTROL: EMBODIMENT

Embodiment: *A mind behind the mirror or a tool to be used? ([Ballsun-Stanton & Torrington, 2025a](#))*

- LIVE DEMO: tokens at perplexity.vercel.app. What your text becomes when the AI “reads” it.
- The context window is what the AI keeps in working memory. Fills up. Things fall out.
- There is no mind in there. The AI is a mirror that reflects your input back, amplified.
- External LOC: AI is a (synthetic) mind with opinions. Internal LOC: It’s a miserable pile of numbers, full of token associations that are shockingly useful.

Where is your judgement right now?

PAUSE.

- Four tensions (Agency, Expertise, Praxis, Embodiment) and the cycle that exercises internal LOC across them.
- When have you taken AI output at face value because it sounded right?
- When have you ignored AI?
- Have you ever decided between different models or services?

What goes inside a good prompt?

ANATOMY OF YOUR PROMPT

Slide 4 was the cycle: what you do across many turns. This is the anatomy of what you write IN one prompt. We will fill it in live for the mindmap question.

Today's task: *Build a Week 7 mindmap node (A History of AI Since the 1950s). The Week 7 lecture is a webpage; we give the AI its URL as context. We fill the prompt live with the class.*

- **Context.** What does the AI need that it can't search for?
- **Register and mechanism.** *How* are we going to accomplish the goal, and with what words?
- **Detail.** Longer prompts beat terse ones.
- **End task.** What should the AI hand back?

Nine prompting strategies: doi.org/10.5281/zenodo.17373463

What did we just make?

DEBRIEF

Workslop is AI content that looks like work but does not meaningfully advance the task ([Niederhoffer et al., 2025](#)). It transfers the work to whoever receives it.

“At the end of it all, slop is a choice. My choice and your choice.” ([Dillard, 2025](#))

- What did the AI ask that you would not have asked yourself?
- Where did the questions feel useful? Where did they feel off?
- How will you carry this back to your own mindmap?

Verify. Edit. Argue.

THREE MOVES OF CRITICAL ENGAGEMENT

“Thou shalt not suffer an error to live.”

- **Verify.** Figure out what the AI’s claims actually are. The AI can help you unpack its own output into discrete checkable claims ([Kudina et al., 2025](#)).
- **Edit.** Edit your prompt sentence by sentence, not the AI’s output. Most “bad AI output” is a prompting problem.
- **Argue.** Ask the AI to challenge your position. The value is in evaluating its counterarguments, not accepting them ([Kudina et al., 2025](#)).

The ability to words good was always the point.

THE TRIUMPH OF THE HUMANITIES

Knowledge has become cheap; judgement has not. ([Jones, 2025](#))

- Figuring out what a text actually says, and articulating your own thinking clearly enough that someone could challenge it, is precisely what determines whether the AI gives you something useful or something dangerously plausible.
- The AI cannot judge whether what it produced is any good. That capacity is what your humanities training has been cultivating, slowly, for years before this lecture ever happened.
- Because the tools will keep changing every few weeks, the only thing that transfers is the skill you came here to develop.

Notes and References

- Ballsun-Stanton, B., & Torrington, J. (2025a). *AI-Locus of Control Continuum* [Graphic]. Zenodo. <https://doi.org/10.5281/ZENODO.17823628>
- Ballsun-Stanton, B., & Torrington, J. (2025b). *Effective AI Prompting Strategies from the Humanities*.
<https://doi.org/10.5281/ZENODO.17373463>
- Ballsun-Stanton, B., & Torrington, J. (2025c, June 3). *Teaching the Unknown: A Pedagogical Framework for Teaching With and About AI*.
<https://zenodo.org/records/15583013>
- Böckeler, B. (2026). *Harness engineering for coding agent users*. martinowler.com. <https://martinfowler.com/articles/harness-engineering.html>
- Darabi, A., Arrington, T. L., & Sayilir, E. (2018). Learning from failure: A meta-analysis of the empirical studies. *Educational Technology Research and Development*, 66(5), 1101–1118. <https://doi.org/10.1007/s11423-018-9579-9>
- Dell’Acqua, F., McFowland, E., Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>
- Dillard, J. E. (2025, October 9). *Slop is a choice*. jdilla.xyz. <https://www.jdilla.xyz/post/slop-is-a-choice>
- Jones, N. (2025, June 2). *I Summarized Mary Meeker’s Incredible 340 Page 2025 AI Trends Deck—Here’s Mary’s Take, My Response, and What You Can Learn* [Substack newsletter]. Nate’s Substack. <https://natesnewsletter.substack.com/p/i-summarized-mary-meekers-incredible>
- Kapur, M. (2008). Productive Failure. *Cognition and Instruction*, 26(3), 379–424. <https://doi.org/10.1080/07370000802212669>
- Kudina, O., Ballsun-Stanton, B., & Alfano, M. (2025). The use of large language models as scaffolds for proleptic reasoning. *Asian Journal of Philosophy*, 4(1), 24. <https://doi.org/10.1007/s44204-025-00247-1>
- Niederhoffer, K., Kellerman, G. R., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (2025, September 22). AI-Generated “Workslop” Is Destroying Productivity. *Harvard Business Review*. <https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity>
- Rotter, J. B. (1966). Generalized expectancies for internal versus external control of reinforcement. *Psychological Monographs: General and Applied*, 80(1), 1–28. <https://doi.org/10.1037/h0092976>
- Schön, D. A. (1994). *The Reflective Practitioner: How Professionals Think in Action*. Taylor & Francis Group.
<http://ebookcentral.proquest.com/lib/mqu/detail.action?docID=4816972>
- Storey, M.-A. (2026, February 9). *How Generative and Agentic AI Shift Concern from Technical Debt to Cognitive Debt*.
<https://margaretstorey.com/blog/2026/02/09/cognitive-debt/>

Final Mindmap Submission.

DUE SUNDAY @ MIDNIGHT THIS WEEK

- Final Mindmap Submission Due Sunday @ Midnight This Week
- We recommend your final entry include commentary on anything you did to finalize the mindmap from your last version(s), rather than specifically integrating Week 11 content into the mindmap