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?

Speaker notes

Budget: ~3 min

Sources:

- Jones (2025) (epigraph; “knowledge cheap, judgement not”; line 185 of the Substack piece)
- Böckeler (2026) (the “Agent = Model + Harness” framing for the layers bullet)

Plants the lecture’s central question. Judgement is the through-line; the question is *who* is holding it. Recap rhyme: “where is your judgement right now?” will return as the consolidation slide and again at close. Forward-link to slide 1 (LOC intro) for the technical name of the question.

Must-say: Read the Jones epigraph; pose the question; gloss the four layers in 30 seconds.

Layer glosses (voiced when bullet 1 lands, not on-slide):

- *Model* is the engine. Tiered: Claude Opus or Sonnet; GPT-5.5 xhigh or instant; Gemini 3.1 Pro or Flash. Same brand, different tier, very different behaviour.
- *Harness* is the prompts and tools wrapping the model so it can execute your prompts. Bockeler: “Agent = Model + Harness.” Vendor sets it up; you can sometimes add to it.
- *Prompt* is the exact words you type plus the content you include or that gets fetched. The layer most under your direct control.
- *Modality* is the interface (web vs desktop app vs command line). Same model, same harness, different defaults.

Discussion question (on-slide bullet 4 IS the question): Hold this; do not press for answers yet. Slide 1 introduces the term for 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?

Speaker notes

Budget: ~5 min

Sources:

- Rotter (1966) (origin of LOC concept)
- Torrington et al. (2025) (longitudinal AI-LOC study; OSF preprint)
- Ballsun-Stanton & Torrington (2025a) (Continuum diagram; the four tensions)

Plants the framework. This is the term-introduction slide; everything from here forward names tensions from this framework.

Must-say: Define LOC in plain English; then the AI version; then the four tensions; then the question (which side am I on?).

Everyday example to voice (anchor before the AI application):

- Internal LOC says “I did well because I studied; I did badly because I did not.”
- External LOC says “I did well because the questions were easy; I did badly because the lecturer hates me.”
- Most people are not at one pole; they drift, depending on the situation.

Forward-link: Slides 2, 3, 5, 6 each take one tension and unpack it.

The Continuum.

AI-LOCUS OF CONTROL

Open the AI-LOC Continuum diagram (full size)

Speaker notes

Budget: ~1 min (visual companion to slide 1)

Sources:

- Ballsun-Stanton & Torrington (2025a) (the diagram itself)

Plants the visual. The four tensions just named — here they are as a single map. Let the diagram do the work; do not narrate every quadrant.

Must-say: “This is the same framework, drawn out. We will revisit each quadrant as we go.”

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.

Speaker notes

Budget: ~10 min

Sources:

- Ballsun-Stanton & Torrington (2025a) (Agency definition on-slide)
- Willison (2023) (calculator-for-words framing; bullet 2)

Plants Agency. First tension named. You are not selecting from a menu; you construct the input.

Must-say: Tokens are how the AI actually reads what you type; we will see this live on slide 6. The “calculator for words” line is the anti-conception: not search, not a friend.

Token gloss (voiced when bullet 1 lands): rentry.co/samplers puts it well: “LLMs have a vocabulary, or a dictionary, of valid tokens.” A token is a dictionary entry the model knows. “The” is one token; “preprocessing” is three. Your typed text gets converted into a sequence of tokens before the AI ever sees it. Slide 6 has the live demo.

Cooking-recipes line: Land this; it is the most relatable example. Push for “have you ever taken the first AI recipe straight”?

Forward-link: slide 6 (Embodiment) demonstrates tokens live.

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.)

Speaker notes

Budget: ~6 min

Sources:

- Ballsun-Stanton & Torrington (2025a) (Expertise definition on-slide)
- Kudina et al. (2025) (drunk tutor metaphor; bullet 1)
- Torrington et al. (2025) (legal-case student quote; bullet 2)

Plants Expertise. Confidence does not track accuracy. The tonal failure mode is the rest of the lecture's reason for existing.

Must-say: Drunk tutor is concrete; the legal-case quote is the killer; bullet 3 reframes what the AI output actually IS (statistical consequence of words, not authority).

GPT-slip anecdote to deploy: I once asked an LLM about "General-Purpose Transformers" because I had the acronym wrong; it cheerfully invented a definition. Whichever model, whichever tier, whichever harness, the confidence does not track accuracy. (One anecdote per slide; this is the one.)

Week-4 classroom line from the longitudinal study (deploy-if a student says "but the AI sounds so sure"): What I told that class was, "Structuring your ideas, it [AI] cannot think. YOU have thoughts and opinions. Fight it." Same instruction here.

Forward-link: slide 9 (Verify-Edit-Argue) is the operational answer to the drunk tutor.

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.

Speaker notes

Budget: ~7 min

Sources:

- Kudina et al. (2025) (productive failure as the AI pedagogy frame)
- Torrington et al. (2025) (bottom quote)
- Schön (1994) (Schön's reflective practitioner; speaker notes only)
- Kapur (2008) and Darabi et al. (2018) (productive failure background; speaker notes only)

Plants the method. Six steps; each is a workshop handle. The cycle is how you exercise internal LOC across the four tensions.

Must-say: Walk through each beat with a concrete example; emphasise that goals precede choice precede expectation; land the bottom quote as evidence the cycle works in practice.

Schön (speaker notes only; do not name to students): Reflective practice; his four diagnostic questions ("Can I solve the problem I have set? Do I like what I get? Have I made the situation coherent? Have I kept inquiry moving?") shape the "evaluate" step. Productive failure (Kapur, Darabi) is the empirical literature on why failure is the method, not the failure mode.

Deploy-if a student says "this seems like a lot of steps": Of course it is a lot of steps. Reading a primary source is also a lot of steps. The point of training is that the steps become automatic; that is what we are practising.

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).

Speaker notes

Budget: ~5 min

Sources:

- Ballsun-Stanton & Torrington (2025a) (Praxis definition)
- Dell'Acqua et al. (2026) (Jagged Frontier; bullet 1; supersedes the 2023 SSRN working paper, formally published Mar 2026 in *Organization Science*)
- Storey (2026) (cognitive debt; bullet 4; optional tute reading)
- Torrington et al. (2025) (bottom student quote)

Plants Praxis. The when-to-engage question. The frontier is real; choosing well is the skill.

Must-say: The Mollick quote on-slide does the heavy lifting; gloss it in plain English. The citation-manager example is the humanities-relevant analogue (you already use Zotero; the AI is one more tool among several, not the only one).

Mollick hedge to voice: Mollick himself has said the frontier is narrowing as models close gaps. Taste is still the differentiator; the frontier is not gone, just further out.

Cognitive debt (Storey) connection to the tute: Tute discussion question 5 is about exactly this; prime them by naming the concept here.

Forward-link: slide 8's WE demonstrates internal Praxis (using AI to augment your mindmap rather than replace your thinking).

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.

Speaker notes

Budget: ~10 min (longest body slide; live demo carries time)

Sources:

- Ballsun-Stanton & Torrington (2025a) (Embodiment definition)
- `perplexity.vercel.app` (token visualiser; live tool)
- Ballsun-Stanton & Torrington (2025b) (the amplifying mirror framing; the poster)

Plants Embodiment. The mirror has no kitten. The numbers are real and useful; the mind is not.

Must-say: Do the live demo; show tokens at perplexity.vercel.app with a sentence the class supplies. Show the context window filling. Land the “no kitten” image as the takeaway.

Conversation-splitting heuristic (voiced when bullet 2 lands): When the context window fills up, the AI starts losing what you told it ten turns ago. The fix is not “type more loudly”; the fix is start a fresh conversation and bring forward only what you need. Strategy 2 from our guide (provide comprehensive context) does not mean “stuff everything in there”; it means “give the AI what it needs to do this turn.”

Workshop slide 19 callback: “Amplifying mirror; no creativity. Reflects the user’s creativity, judgement, discernment. There is no real kitten in the mirror.” This is the image; it is yours from 2023.

Forward-link: slide 9 (Verify-Edit-Argue) is what you do because the mirror does not know if its reflection is true.

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?

Speaker notes

Budget: ~4 min

Sources: None on-slide. This is a recap, not a content slide.

Plants the consolidation. Mid-deck breath; rhymes with week 7's "where is the intelligence right now?" recap.

Must-say: Read each question aloud, pause, take answers. Resist filling silence. Three discussion questions is a lot; expect this to run to 5 min if the room engages.

Tension unpacks (voiced when bullet 1 lands, one Continuum question per tension):

- Agency: does the AI have agency, or do you?
- Expertise: better than you, or input to your judgement?
- Embodiment: a mind, or a tool?
- Praxis: automation, or augmentation?

Deploy-if the room is silent: Pick one question; pick one student you know will engage; ask them directly. The pause is the engagement; do not rescue it.

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

Speaker notes

Budget: ~12 min (build + run live; workslop + debrief on the next slide)

Sources:

- Ballsun-Stanton & Torrington (2025b) (the nine strategies; our poster)

Plants the WE. The single worked example for the lecture; collapses prior WE0/WE1 into one demonstration on a task students are doing this week.

Must-say: Read the task; build the prompt live with class input (transcribe, do not pre-write); run it on Week 7's URL.

Live build sequence (informal; let class drive the order):

- **Context** first: what does the AI need to know about Week 7 that it cannot search for? (Answer: nothing if we give it the URL.)
- **Register and mechanism** second: do we want it to ask us questions one at a time, or produce something? What words convey that?
- **Detail** third: how much do we tell it about what a mindmap node is?
- **End task** last: what is the deliverable? A list of concepts? A description? Both?

Deploy-if a student asks "what if I just want the AI to do it": That is the external Praxis pole; that is the cognitive-debt move (slide 5). The mindmap rubric will mark you down because the dialogue is missing.

Forward-link: slide 8b carries workslop + debrief.

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?

Speaker notes

Budget: ~8 min

Sources:

- Niederhoffer et al. (2025) (workslop definition; HBR Sep 2025)
- Dillard (2025) (closing quote)

Plants the debrief. Post-run reflection; names the failure mode (workslop) the live run was trying to avoid; bridges to slide 9’s three moves.

Must-say: Read the workslop definition; land Dillard. Then take each debrief question in turn. Resist filling silence. Bullet 3 (how will you carry this back) is the bridge to the mindmap submission this Sunday.

Workslop tie-in: The HBR paper is recent (Sep 2025) and corporate-flavoured, but the definition is clean. When you share AI output you have not validated, you have shifted the work to whoever has to deal with it; that is the cost of workslop. Connects to the mindmap rubric: “show your work” exists because the alternative is workslop. JD Illa’s “slop is a choice” line names the broader pattern; the HBR paper names the work-context variant.

Forward-link: slide 9 (Verify. Edit. Argue.) is what you do with what you just produced.

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).

Speaker notes

Budget: ~10 min

Sources:

- Kudina et al. (2025) (verify, edit, argue framings; proleptic reasoning; “value is in the evaluation, not the counterarguments”)
- Ballsun-Stanton & Torrington (2025c) (grimoire annotation as a deploy-if tactic for “edit”)

Plants critical engagement. Three moves; absorbs the earlier-deck plans for grimoire annotation and proleptic reasoning into Verify-Edit-Argue.

Must-say: Each move has a concrete demonstration; do not abstract. Verify: ask the AI to list its claims as bullets. Edit: rewrite the prompt, not the output. Argue: ask the AI to attack your position; evaluate its attacks.

AI-can-list-claims tactic (under Verify): A prompt like “list the discrete claims in your previous response as a bulleted list, marking each as factual claim, opinion, or hedge” makes verification mechanical. The AI is bad at knowing what is true; it is reasonable at categorising its own output if you ask.

Proleptic reasoning (voiced under Argue; do not name the term on-slide): Anticipating counterarguments before they are made; the AI can simulate a hostile reader. The value is in your evaluation of the simulation, not the simulation itself.

Deploy-if a student asks for a concrete editing tactic: Grimoire annotation (BBS/Torrington 2025; pink/blue highlighting in Figure 1 of *Teaching the Unknown*). One named tactic; not a parade.

Forward-link: slide 10 closes on what these three moves cultivate (taste and judgement, the durable edge).

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.

Speaker notes

Budget: ~4 min

Sources:

- Jones (2025) (Durable Human Edge framing; the epigraph callback to slide 0)

Plants the close. Triumph payoff; the title returns; Jones bookends the lecture.

Must-say: The Jones epigraph reappears as the through-line; the three bullets land in order. Do not rush; this is the rhetorical bookend.

Tute hand-off: “Specify your own program” tute activity directly uses the prompt anatomy you just practised on slide 8. The mindmap submission this Sunday is where you apply it.

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*. martinfowler.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., Krayer, 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

Speaker notes

Budget: ~2 min

Must-say: Read the notice. Take questions. Done.