

Lecture IX — Judgment and Presenting

AI-Assisted Programming for PhD Researchers

Dr. Tobias Vlček
Helmut Schmidt University — September 2026

When Not to Use AI

The Question Is Not “Can It?”

- It usually *can* produce something — that was never the question
- The question is whether **you** should delegate this
- Now, in a thesis you will have to defend
- “It compiles” and “I can stand behind it” are different bars

Do Not Delegate

- What you are here to **learn** — the core methods of your field
- **Architectural decisions** you will live with for years
- Anything you **cannot verify** yourself
- **Statistics** you do not actually understand

The Verification Boundary

- Usable delegation ends where your ability to **check** ends
- Self-perception is a bad meter for this

...

- METR’s developers *felt* ~20% faster while being **19% slower** [1]
- Growing that boundary outward **is** your development as a researcher — the goal is calibrated trust

Keeping Your Edge

- Deliberately code **without** AI sometimes — small katas, first drafts of core logic
- Make **explain-back** a habit: say what the code does before you accept it
- Confidence in the tool predicts *less* critical thinking [2]
- The RCT, one last time: **how** you use it decides what you keep [3]

Your Own Guardrails

- When the goal is learning, ask the agent for **hints and next steps**, not full solutions
- That is the exact design that erased the learning harm in the PNAS tutoring RCT [4]
- The effect is specific to that guarded-tutor study — treat it as a pattern to copy, not a guarantee
- After this course, **you** are your own guardrail

Discussion

Note

Two questions for the room:

- Which task this week would you **not** delegate again?
- Where did you accept something you couldn't explain?

Presenting Your Project

The Shape of a Good 12 Minutes

- **Problem** — what and why (2 min)
- **Live demo** — show it working (4 min)
- **Process** — how you worked with the agent (3 min)
- **Challenges** — what broke, what you changed (2 min)
- **Insights** — what the room should take away (1 min)

Mirrors the four points you are graded on.

Demo Risk Management

Tip

Live demos fail. Have a fallback:

- Record a 60-second screen capture **now**, while it works.
- Test on the projector's resolution before you present.

Slides, Fast

- You may of course use AI for the slides too
- Same rules apply: **every claim on a slide is yours**
- Keep it lean — aim for ~8 slides
- Talk to the demo, not to a wall of text

Reflection Beats Perfection

- An honest “here is where it went wrong and what I changed”...
- ...scores better than a polished façade
- It also teaches the room more than a flawless run
- Your **judgment** is the thing on display, not the tool's output

Continue Your Journey

After the Course

- Keep the three habits: the **loop**, the **checks**, the **review**
- This website stays up — come back to it
- Go deeper: [course literature and references](#)

- Details for the final talk: [Presenting Your Project](#)
- Questions later? tobiasvlcek.com

Bibliography

- [1] J. Becker, N. Rush, E. Barnes, and D. Rein, “Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.” Accessed: July 09, 2026. [Online]. Available: <https://arxiv.org/abs/2507.09089>
- [2] H.-P. (. Lee *et al.*, “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Yokohama Japan: ACM, Apr. 2025, pp. 1–22. doi: [10.1145/3706598.3713778](https://doi.org/10.1145/3706598.3713778).
- [3] J. H. Shen and A. Tamkin, “How AI Impacts Skill Formation.” Accessed: July 09, 2026. [Online]. Available: <https://arxiv.org/abs/2601.20245>
- [4] H. Bastani, O. Bastani, A. Sungu, H. Ge, Ö. Kabakçı, and R. Mariman, “Generative AI without guardrails can harm learning: Evidence from high school mathematics,” *Proceedings of the National Academy of Sciences*, vol. 122, no. 26, p. e2422633122, July 2025, doi: [10.1073/pnas.2422633122](https://doi.org/10.1073/pnas.2422633122).