Co-Creating a Class AI Policy

Tim Rodgers, Sara Harris

1. What motivated you to invite student participation to develop your class generative AI policy?

ENVR 200 “Introduction to Environmental Science” is a second year, writing and communications intensive class which counts towards our students’ communications credit requirement for the Faculty of Science. As part of a broader effort supported by a Faculty of Science AI Curriculum Initiative grant awarded to Tara Ivanochko, Sara Harris and Jason Young, with implementation support from Warren Code at Skylight, we sought to introduce students to thoughtful and responsible uses of generative AI (GenAI). The grant identified ENVR 200 as a natural venue for introducing content around the responsible and ethical use of AI because it is the first core course students take after joining the Environmental Science Specialization and therefore helps establish program-wide norms and expectations.

Pedagogically, initial research suggests that AI appears to be harmful to learning when it is relied on for easy answers that replace thinking, but that when appropriately scaffolded AI can enhance learning, especially once a learner has a baseline for a skill (Bastani et al., 2025; Dong, 2026; Han et al., 2025; Kosmyna et al., 2026). Together, this indicates that using AI to enhance student learning requires students to be deliberate and intentional in their AI use, which in turn requires self-motivation.

This meant that our goal was to develop a learning activity that would encourage students to critically reflect on how their use of AI impacts their learning. As instructors, a key tool that we have to shape student use of AI is the course AI policy. AI policy hasn’t yet converged on an optimal form – many people are experimenting with AI in different ways, including UBC students who have shown in student surveys that they have diverse opinions on what constitutes an acceptable use of AI (CTLT, 2026). This lack of convergence means that while instructors have freedom to determine their own AI policies, they may simultaneously may not feel confident as to what the best AI policy is. Compounding this problem, the “best” AI policy also varies greatly between different classroom contexts and even within a class between different assignments, and enforcement of AI policy is notoriously difficult with tools like AI checkers known to be biased and imprecise, making them inappropriate for individual decisions (Erol et al., 2025; Gotoman et al., 2025).

2. What learning objectives were you trying to address?

We therefore decided to create an activity where our students helped to co-create the ENVR 200 class AI policy. We felt that involving students in a decision that directly affects

their learning experience could increase buy-in by providing greater agency and ownership over the process (Ryan & Deci, 2000). We also hoped to draw on students' expertise in their own AI use to find ideas or mechanisms of an AI policy that could enhance student learning. Our primary learning objectives were that, by the end of this activity, students should be able to:

- Reflect on how generative AI tools impact your learning, such that you can have a conversation with a classmate about pros and cons of GenAI use.
- Discriminate between uses of generative AI that will enhance your learning and those that will not by deciding on an appropriate class policy for the use of generative AI tools.
- Make reasoned, ethical decisions about AI use in one's own oral communication and writing

This approach builds on the work of Zonjić, (2024), who initiated class discussions on the use of AI and used those discussions to inform her AI policy.

3. What did your activity look like, and how did it incorporate GenAI?

We conducted our AI policy co-creation activity during portions of two of our 3-hour class periods (Lecture 1 and Lecture 2), preceded by a pre-class survey. For the main drafting and ratification activities we broke the class into groups and directed them to a google-docs template with tabs for each group, in which each group drafted their own policy. All materials are available from our activity repository.

Day 0: Pre-Class Survey

We started with a pre-class survey asking students how they have used GenAI in the past, and whether they felt that their use of AI helped or hindered their learning. The goals of the pre-term survey were to prime the students to think critically about their AI usage, provide information that we could share within the class on how students use AI, and provide us as the instructors with information on student attitudes towards GenAI that could inform the policy. For example, if the students overwhelmingly said that they felt AI was hindering their learning, we could point to that result to justify implementation of a more restrictive policy (or vice-versa).

Day 1: Co-Drafting the Class AI Policy

We introduced the AI policy co-creation activity during the first lecture of the term after going over the course deliverables, so that students would understand course assignment expectations and how they might use AI for those assignments.

Next, we contextualized the activity by going over student responses to the pre-class survey and followed up with a brief literature review on GenAI use in learning. Here we explained that this is an emerging field and there aren't any firm answers on how GenAI impacts learning, but that preliminary evidence indicates that using AI to replace learning (e.g. by asking "what is the answer") seems to lead to worse outcomes than using AI to support learning (e.g. "I don't understand X concept, can you explain this in different words"). We also emphasized that the evidence suggests that AI can be more effectively used to support learning once you have a baseline skill and highlighted that although AI can make students feel like they are learning quickly (Zohar et al., 2026), overall learning can be enhanced by struggling through the content and experiencing "desirable difficulties" that help with long-term learning and retention (Bjork & Bjork, 2011).

We then discussed the components of an AI policy (modified from (Stanford Center for Teaching and Learning., n.d.), drawing on our previous contextualization to show that the goal of the policy is to support learning and what that looks like is still unclear and may vary for different learning activities and skill levels, even within the same class. We summarized this with the following policy outline:

- Nature of policy (e.g. Unrestricted, partially restricted, total ban)
- Applicable conditions or contexts, including:
    1. Allowable uses (e.g. editing, spellcheck, background summarizing, etc.)
    2. Restrictions (e.g. may not use for Assignment XX, only with permission, not for first draft)
- Process for compliance, including:
    1. Attribution guidelines
    2. Compliance auditing
- Context and rationale for policy
- Consequences for non-compliance
- Support resources for students

For transparency we also introduced guard-rails from the instructional team based on our course needs and AI positionality stating that the final policy must ensure that students:

- Generate their own initial ideas, with writing that reflects their understanding
- Write all first drafts to maintain ownership and understanding of their work
- Thoroughly document and disclose all use of AI
- Do not directly submit AI-generated material for any assessed work in this class (unless stated directly in the assignment)

We then had students draft AI policies in small groups by going through the AI policy framework step-by-step using an outline we provided in a google-docs template.

1. First, we had the students determine the overall nature of their AI policy and what the conditions would be for the different assessment types in our course.
    a. For our class we broke this down into personal reflections (our class has weekly journals), written assessments, and presentations, and asked the students to give examples of allowable uses and explicit restrictions for each.
2. Next, we asked students to draft what we called "Processes for Compliance". Here we included guidelines around AI attribution statements, along with "compliance auditing" mechanisms that could both help the student provide evidence that they hadn't used AI inappropriately and help the instructor ensure that AI was being used as per the policy.
3. Finally, we asked students to review their draft policies and to:
    a. Draft a rationale statement outlining why they chose their policy, reflecting on how this policy will help them learn, who is helped or harmed by the policy, and whether their policy meets their idea of what an AI policy should do.
    b. Reflect on how it would feel to complete one of the course assessments following this policy
    c. Iterate as needed until they find a policy that aligns with how they think AI should be used to support learning.
    d. As their exit ticket we had students reflect on two questions:
        - What are you taking away from our discussion about AI today?
        - What did you think about, find surprising, or learn?

Initially, we included drafting the rationale as a first step, but we found that the rationale drafted by the students at the beginning of the assessment was superficial. Moving forward, we suggest having learners draft the rationale after they have looked at the different assessment types (as presented here) so that the students can think about it think about the rationale in context. Overall, the critical piece is to allow enough time for the students to review their policy to ensure that students consider alignment between the policy they have created and their stated rationale for that policy.

Day 2: AI Policy Discussion, Debate & Votes

The instructors reviewed each group's draft policies between lectures, first summarizing the major themes and then using that to draft an overall policy. We implemented any policy that fell within our guard rails which had the support of a clear majority of student opinion. A good example here was to allow the use of GenAI grammar and spelling tools (7/9 groups supported, ~2 called for a total ban), and not to allow using GenAI tools for creating a first

draft (7/9 groups explicitly mentioned this and none called for allowing GenAI use in a first draft).

For areas without clear majorities, we followed two approaches. For some, use-cases we made a decision as the teaching team. For example, we banned using GenAI for initial idea generation and brainstorming as we felt that it went against our learning goals, although the class was split on this use case(3 allowed, 3 disallowed, and 3 did not mention this policy). For the remainder of items without clear consensus where we felt that using GenAI would not impede our learning goals, we put the question to a vote. We put five items to the vote:

1. Using AI to search the literature and find references
2. Using AI to summarize publications
3. Using AI to generate images for figures
4. Implementing a "version history" policy, where students would need to complete their assignments in a program that tracked version history over time. The idea here was that the teaching team could then look at the version history for evidence of large copy-paste events that would provide one source of evidence for inappropriate AI use.
5. Requiring the submission of a preliminary "human draft" that was created without AI for assignments where AI was used to help.

The "human draft" was an idea that emerged from the review of student policies. We felt that requiring this draft could encourage the students to try creating their deliverables by themselves first and thereby encourage a process where they experienced some useful friction and tried to develop skills like connecting ideas by themselves before using the AI to polish their work and identify if they had missed anything substantial.

During Lecture 2 we reviewed the overall draft policy and had students vote on the outstanding issues via an anonymous online platform. We allowed students to debate each vote, following a public forum format where we called for speakers to motivate either for or against each proposal as it was considered. In the future, we would structure the voting activity to give students more space to consider the proposal individually and in small groups, and we would provide students with an anonymous discussion board that students could use live to contribute without needing to speak in front of the class (the template slides provided propose a structure for this giving 6 minutes per vote). Alternatively, in-class voting during Lecture 2 could be replaced by an anonymous vote and discussion-board on Canvas or another application outside of class time.

4. What were your takeaways from your GenAI activity?

As has been noted elsewhere (CTLT, 2026), in our pre-class survey we found that our students had a diverse range of opinions on the use of AI in teaching and learning. Reading

the draft AI policies from the student groups gave us the opinion that this activity successfully encouraged students to think-through the effects of using AI, and hopefully gave them some additional tools for deciding when the use of AI is appropriate to support their learning vs when it might not be.

We were surprised when the majority of students chose the more restrictive policy for each vote (Table 1), although the margins were extremely small for some votes, (e.g. the more restrictive side won by one vote on the question around using AI for image generation).

Table 1: Results of class votes on the five proposed policies (see the google doc for full wording). For votes 1-3 the policies were worded as "Students may/may not use generative AI for [XX]" (e.g. vote 1 XX = "summarizing literature), while for votes 4-5 the policies were worded as "the class will follow the [XX] policy outlined below" (e.g. vote 3 XX = "version history"). The winning response is in bold.

Vote #	Class Votes	Yes	No	Winning Share
1	Searching literature	19	**23**	55%
2	Summarizing Publications	11	**32**	74%
3	Image generation	21	**22**	51%
4	Version history	**24**	18	57%
5	"human draft"	**23**	19	55%

We plan to follow-up on this activity with a survey closer to the end of term where we will ask the students how they felt about the co-created AI policy itself and the process of creating it. Example questions could include asking if the policy co-creation activity changed their feelings of autonomy around AI use for our class, and if the co-creation process changed how likely they were to follow the policy, particularly in cases where they were on the "losing" side of a narrow split.

5. Would you run this activity again? How could other instructors implement this activity?

Overall, this activity took ~2 hours across two lectures. This amount of time made sense for our class as we were piloting the activity, the class is focused on writing and communications, the class aims to set a foundation for student expectations in our program, and the class has a 3 hour timeslot once per week. This is likely too long for most classes.

To shorten this activity, we recommend either providing example text for students to pick between (e.g. "students may / may not use generative AI for XX task in YY assignment") in the template and/or limiting the amount of discussion at each stage depending on the time available. The template we have provided incorporated some of these ideas and would take ~1.5 hours, spread across two lectures. As discussed above, the Lecture 2 activity could be done asynchronously to reduce in-class time.

As one example, Prof. Amanda Giang successfully adapted this activity for her CEEN 525, Energy Policy course in UBC's Clean Energy Engineering Master's of Engineering Leadership program, taking ~30 minutes of class time. CEEN 525 is a small cohort graduate seminar of students with extensive professional experience. For this class, Prof. Giang discussed how her students used AI in their professional lives and in their studies, and how appropriate GenAI use might differ for learning vs professional activities. Then, she broke down the course learning activities into categories (e.g., completing course readings) and defined guard-rails, such as for an in-class reading reflection assignment where Prof. Giang decided not to allow AI. For each type of assessment, the CEEN 525 class took a straw poll and moved on quickly if there was consensus, and discussed in more depth what AI use was appropriate for assessment types where opinions differed. After this lecture, Prof. Giang summarized the agreed-upon policy and posted it on Canvas as the official AI policy.

Prof. Giang found that students saw distinctions between using GenAI for learning and using it for professional tasks and the students felt that more restrictions were appropriate in a learning context than at work. She also found that students for whom English is an additional language highlighted how AI could enhance their learning in reading and communication heavy classes, something we also found in ENVR 200. Overall, this emphasized the need for a context-aware approach to AI policy that considers learner make-up alongside course and program goals.

We encourage anyone interested in running this activity to download our lecture materials from our activity repository, and please reach out to us if you have any questions or would like help implementing it!

Acknowledgments

Many thanks to Tara Ivanochko and Jason Young for working with SH in conceptualizing the initial grant; to Amanda Giang for sharing her experiences, to Warren Code for the support from Skylight, to Sarah Bean Sherman for her thoughtful input and comments, and to Kate Moloney for assistance evaluating and summarizing the AI policy workshop results.

AI Disclosure Statement

We used Microsoft 365 CoPilot to edit portions of this document at the sentence and paragraph level. All AI suggestions were reviewed by the authors before being incorporated into this document.

Works Cited

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. Proceedings of the National Academy of Sciences, 122(26), e2422633122. https://doi.org/10.1073/pnas.2422633122

Bjork, E. L., & Bjork, R. A. (2011). Making Things Hard on Yourself, But in a Good Way: Creating Desirable Difficulties to Enhance Learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), Psychology and the real world: Essays illustrating fundamental contributions to society. Worth Publishers.

CTLT. (2026). How UBC Students are Using Generative AI: A 2026 Update. AI In Teaching and Learning. https://ai.ctlt.ubc.ca/how-ubc-students-are-using-generative-ai-a-2026-update/

Dong, Y. (2026). Generative AI technologies and educational outcomes: A comprehensive meta-analysis comparing traditional and AI-driven approaches. Humanities and Social Sciences Communications, 13(1), 559. https://doi.org/10.1057/s41599-026-06903-y

Erol, G., Ergen, A., Gülşen Erol, B., Kaya Ergen, Ş., Bora, T. S., Çölgeçen, A. D., Araz, B., Şahin, C., Bostancı, G., Kılıç, İ., Macit, Z. B., Sevgi, U. T., & Güngör, A. (2025). Can we trust academic AI detective? Accuracy and limitations of AI-output detectors. Acta Neurochirurgica, 167(1), 214. https://doi.org/10.1007/s00701-025-06622-4

Gotoman, J. E. J., Luna, H. L. T., Sangria, J. C. S., Santiago Jr., C. S., & Barbuco, D. D. (2025). Accuracy and Reliability of AI-Generated Text Detection Tools: A Literature Review.

American Journal of IR 4.0 and Beyond, 4(1), 1–9. https://doi.org/10.54536/ajirb.v4i1.3795

Han, X., Peng, H., & Liu, M. (2025). The impact of GenAI on learning outcomes: A systematic review and meta-analysis of experimental studies. Educational Research Review, 48, 100714. https://doi.org/10.1016/j.edurev.2025.100714

Kosmyna, N., Hauptmann, E., Yuan, Y. T., & Situ, J. (2026). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task△.

Ryan, R. M., & Deci, E. L. (2000). Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. Contemporary Educational Psychology, 25(1), 54–67. https://doi.org/10.1006/ceps.1999.1020

Stanford Center for Teaching and Learning. (n.d.). Creating an AI Course Policy Workshop Kit | Teaching Commons. Retrieved October 2, 2026, from https://teachingcommons.stanford.edu/professional-development/workshops-programs/do-it-yourself-workshop-kits/creating-ai-course-policy

Zohar, E., Bloom, P., & Inzlicht, M. (2026). Against frictionless AI. Communications Psychology, 4(1), 39. https://doi.org/10.1038/s44271-026-00402-1

Zonjić, M. (2024). Need a policy for using ChatGPT in the classroom? Try asking students. Nature, d41586-024-01691–01694. https://doi.org/10.1038/d41586-024-01691-4