

# Scientific Judgment in the Third Age

Peter Huybers, Chair, Department of Earth and Planetary Sciences  
September 5, 2026

## What has changed

David Deming told the Class of 2030 at [Convocation on Tuesday](#) that universities have had two ages and are entering a third. In the first, Harvard was the place that held the books. Knowledge was scarce and lived on shelves, and a college was where you went to read it. When print became cheap and knowledge specialized, the university became the place that held the experts. That is the university most of us have known. AI now makes expertise cheap and personal. A model answers your question immediately in as much detail as you want. Deming's claim is that the third age has begun but is yet to be built. Our faculty met this week to discuss AI in teaching, and much of the conversation was about what this third age should contain.

Two committee reports frame the problem. [MIT's committee](#), reporting in August, found that AI already produces credible answers to nearly every written assignment in its undergraduate curriculum. An FAS roadmap written in February bracketed the future between an "AI fizzle," in which capabilities plateau, and an "AI boom," in which they keep compounding, and judged the boom to be the case for which the faculty is least prepared. It cautioned readers whose reaction is "but AI can't do X." Jeremy Bloxham made a similar point at our meeting: the technology is moving fast enough that what we know today could be wrong in six months. Every strength and limitation described below is an observation from the past year, and the balance may look different in 2027. For this reason, MIT's committee made humility its first principle.

## What we teach

Ann Pearson put the emphasis on judgment. Judgment means having the experience and background to know which answers are right, and, beyond that, knowing what constitutes good work. A question can have a correct answer and still be a dull question, and a calculation can be right and still miss the interesting part of the problem. A student who can tell the difference has learned something no answer key contains. The FAS roadmap makes a similar observation that knowledge work is shifting away from production and toward evaluation, judgment, and decision-making, and that coursework should mirror the shift.

Scientific judgment also includes knowing the uncertainties. We are used to people being wrong. Over years we learn whom to trust and by how much, and we calibrate accordingly. We do not yet have that track record with AI, and its failure modes are unfamiliar. It is fluent, confident, and wrong in ways we do not always expect. The question raised at the meeting was how to teach students to calibrate a source that is right most of the time and wrong without warning, and how to cultivate a healthy working relationship with it. A scientific answer is to treat AI output as a hypothesis that should be checked against observations and established theory.

Checking an answer, and knowing what to ask next, depend on understanding the fundamentals of a field. Daniel Jacob, who is updating his textbook on atmospheric chemistry, observed that the value of writing a textbook now lies in giving students a clear guide to those fundamentals. The FAS roadmap asks us to consider which tasks students can now automate and forget, and which skills they should still master. Gaining that mastery seems a prerequisite for using AI as a productive tool.

Fiamma Straneo reminded us that the point of learning is more than a student producing the right answer from their own knowledge. There must be room for exploration, inspiration, and fun. This is something we can foster in classrooms and research through team-based and project-based work, where the learning is in how you get there.

## **How judgment forms**

Mastery and judgment have traditionally come from working through hard problems, and Roger Fu asked how students should gain them when answers are readily served up by AI. There was general agreement that the traditional approach worked for those of us around the table, but also interest in whether other approaches are now accessible. Perhaps it is possible to work through research-level projects in new and exciting ways that also develop good scientific judgment.

For all its potential, AI may already be impeding student progress. Several of us saw a pattern last year in which students did well on homework but much worse on closed-book, closed-computer tests than students in previous years. When the prompting and collaboration that AI supplies were removed, students found it hard to make progress unsupported. This supports the notion that an AI partner on homework can displace the productive struggle that mastery requires. The homework was still measuring something, but how it was approached may have been poor preparation for the test.

Frank Keutsch and Zhiming Kuang both pointed to individual oral exams as a good way to assess what a student knows. Ten minutes of conversation shows whether a student can reason without help, and it teaches at the same time. The Bok Center's guidance on [making assignments AI-resilient](#) points the same way in recommending assessment of key competencies in supervised settings and following take-home work with short conversations in which students explain or extend their work.

Steve Wofsy, who proposed that we discuss the MIT report, was drawn to its emphasis on the human component of learning, saying this may be the crux of what a university should offer in the third age. Deming's own answer was that universities should become more communal. An EPS education has some of this already built in. Reading groups and field trips are communal by nature, and laboratory and field measurements are among the more AI-durable parts of science.

The same considerations apply to research training. The FAS roadmap cautions that some faculty may come to see graduate students as a labor supply that AI can replace. MIT's

committee heard the same about undergraduate researchers. Both reports answer that students are here to learn, and that research on a campus is an apprenticeship as well as a means of producing results. Faculty around the table also added that students bring their own insights and inspiration, develop and make the observations that underpin much of the science we produce, help teach courses, and go on to play important roles within and outside of academia.

None of this is settled. The discussion pointed toward stating in every course what use of AI is permitted and why, shifting assessment toward conversation and work done in the room, and strengthening the communities within which we work and learn. How we take these directions up is for the department and our broader community to work out over the coming months, and I hope you will weigh in. Whatever we decide should be revisited by spring, when things will look different again.

## Acknowledgements

I used Claude Fable 5.1 to provide feedback and suggested edits in writing this piece. Useful feedback was provided by faculty quoted in this document.

## Sources

- David Deming, Convocation address to the Class of 2030, September 1, 2026: [Crimson coverage](#).
- MIT, *Report of the Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training*, August 13, 2026: [PDF](#); [project page](#).
- Harvard FAS, *FAS and Artificial Intelligence: A Roadmap for 2026–2030* (Claybaugh, Dench, Hall, Nesson, Pierce, Stubbs, Thornber, Jelinkova, Juraschek), February 2026, internal document. Public statement: [Dean Claybaugh on AI](#); [OUE generative AI guidance and sample syllabus policies](#).
- Derek Bok Center: [Teaching in the Age of AI](#); [Making Assignments AI-Resilient](#); [AI Syllabus Policies](#); [Guidance for Faculty on Addressing AI-Related Academic Integrity Issues](#).