

The argument, in short

We are at the dawn of AGI, and essentially every institution is answering it with tool training — much of it authored by the vendors whose tokens the tools consume. That is a category error. "What can the machine do?" is a technology question with a shelf life of about a quarter. "What am I for, once it can do what I was trained to do?" is a humanities question with a shelf life of a career. Berkeley is the right place to ask the second one, because it is the only kind of institution whose entire purpose is allocating opportunity to people who weren't born holding it — and AGI is, before anything else, a re-allocation event.

Why Berkeley specifically — verified

It's a public allocation machine at scale. ~46,000 students. Berkeley doesn't exist to concentrate advantage; it exists to distribute it. If AGI resets who gets access to capability, a public university is the institution with standing to care.

Its AI posture is already distinctive, and it's in the contracts. The BAIR Open Research Commons brings in Apple, Google, Meta, NVIDIA, Samsung, Sony, Bosch and GM on Berkeley's terms: non-exclusive, open publication expected, and Berkeley explicitly discourages publication restrictions or exclusive commercial licenses — with students leading project design. That is a governance stance, not a slogan. CHAI exists to make AI "provably beneficial." CDSS is Berkeley's first new college in 50+ years, built on a \$252M anonymous gift — the largest in its history — and 1 in 5 undergraduates now take a data science course annually.

And there's a conspicuous gap this course fills. UC adopted eight Responsible AI Principles in 2021 — not one addresses labor, employment, or displacement. Berkeley's own Academic Senate guidance (Aug 2025) recommends rather than requires, and leaves judgment with the instructor. Berkeley has deliberately left the "what does this mean for me" question open. Nobody is currently teaching it.

Why a sociologist is instructor of record

This is the strongest and least obvious part of the pitch. Petersen's life's work is exactly this question, thirty years early. His framework — allocative vs. within-job vs. valiative discrimination — locates unfairness not in bad actors but in the opportunity structure: it survives wherever an act is hard to document, ambiguous when documented, and has no complainant. He identified hiring as "the point where discrimination is most feasible."

That is precisely where AI now sits. Pairing him with Bayen (engineering) is the coin made faculty: the person who knows how opportunity gets allocated, and the person who knows how the system doing the allocating is built.

Two more facts make this concrete: Sameer Srivastava — who chairs Petersen's own group at Haas — published on generative AI in organizational selection in Sept 2025; and Solène

Delecourt, same group, showed in Nature (Oct 2025) that ChatGPT rates older male applicants higher when evaluating résumés. The subject is already alive in his corridor.