End of semester:

- Ethics day 1 (Monday)
- Ethics day 2 (**Today**)
- Ethics day 3 (Friday)
- Presentations day 1 (next week Monday)
- Presentations day 2 (next week Wednesday)
- Review for the final exam (next week Friday)
- Final exam (Wednesday, May 7, 10:30-12:30)

Today:

- "Dual use of artificial-intelligence-powered drug discovery"
- "Amazon scraps secret AI recruiting tool that showed bias against women"
- "AI-driven dermatology could leave dark-sknned patients behind"

1. In “Dual use of artificial-intelligence-powered drug discovery” (Urbina et al) and the related reporting, one notable aspect is how easy it is for someone with a little bit of biochemistry knowledge, a laptop, and some open-source software to design novel toxins. Compared to other historical technological advances with lethal/harmful uses, does the ease-of-use and accessibility of machine-learning toxin discovery represent something novel? Or, can we reason about the ethics of machine learning in the same way as other technological advances?

2. “Amazon scraps secret AI recruiting tool that showed bias against women” (Dastin)... but humans are also biased. Is there any hope for developing tools for screening job applicants—or similar tasks—that are *less biased* than humans and hence reduce unfair discrimination?

3. “AI-driven dermatology could leave dark-skinned patients behind” (Lashbrook) ends with researcher commenting, “AI isn’t bad; quite the opposite. I just think it should be inclusive.” It is that simple? Is better training data all it takes for more accurate and accessible diagnostic tools?

Coming up:

Due Thurs, Apr 24:
Read and respond to ethics readings on AI and human learning (See Canvas)

Apr 28 and 30: *Project presentations*

Due Apr 30: *Final submission of notebooks*

(That's all.)