

Tech Giants, Gorging on AI Professors Is Bad for You

If industry keeps hiring the cutting-edge scholars, who will train the next generation of innovators in artificial intelligence?

By [Ariel Proccaccia](#)

January 7, 2019, 6:00 AM EST



Eat too much and there won't be grass for anyone. *Photographer: William West/AFP/Getty Images*

In an essay written in 1833, the British economist William Forster Lloyd made a profound observation using the example of cattle grazing. Lloyd described a hypothetical scenario involving herders who share a pasture, and individually decide how many of their animals would graze there. If few herders exercised restraint, overgrazing would occur, reducing the pasture's future usefulness and eventually hurting everybody.

The sinister beauty of this example is that the rational course of action is to behave selfishly. That's because the selfish herder's cattle would be able to gorge on the pasture as long as considerate herders held their animals back. And there'd be no short-term benefit to selfless behavior: If other herders are selfish, overgrazing of the common agricultural land would occur anyway. This "tragedy of the commons," in Lloyd's words, is a prominent type of social dilemma, where unchecked self-interest on the part of individuals leads to poor outcomes for their group.

Now the tragedy of the commons is playing out in the field of artificial intelligence, with companies as the herders and professors as the grass. ¹

As AI frenzy engulfs the technology and financial sectors, attempts to hire AI experts from the nation's top universities have skyrocketed. Universities cannot begin to compete with the seven-figure salaries that are routinely offered by major companies and even some nonprofits.

And the money isn't everything. For some professors, part of the appeal comes from access to treasure troves of data and awesome computational power – two of the engines that drive applied AI research. Others find in industry a source of compelling, large-scale problems.

The result can be summed up in one word: overgrazing. The School of Engineering and Computer Science at the University of Washington, one of the finest in the U.S., has been among the hardest hit. By my count, eight out of 11 tenured faculty members working in robotics, machine learning and natural-language processing are currently on leave, or spending at least 50 percent of their time at Amazon, Facebook Inc., Apple Inc., Nvidia Corp., D.E. Shaw & Co. and the Allen Institute for AI. ²

Stanford University and Carnegie Mellon University (where I am a faculty member) – both world-leading centers of AI research and education – have recently seen the departures of a Who's Who of AI researchers. More generally, few academic AI groups have escaped unscathed, and more than a few have suffered debilitating losses to industry.

The tragedy is that this state of affairs sabotages the long-term interests of the very companies responsible for it. To understand why, it is important to appreciate the most distinctive aspect of a professor's job: the training of Ph.D. students, an excruciatingly slow and intellectually transformative apprenticeship process. ³

Most of the recent breakthroughs in AI that are driving commercial applications originate in the research of Ph.D. students (AlexNet, generative adversarial networks, and Libratus, just to name a few). When these Ph.D. researchers graduate, they are heavily recruited by companies, which recognize that their importance to the progress of AI in industry is only matched by their scarcity. Still, some of the most formidable Ph.D. graduates remain in academic life, becoming professors who train more Ph.D. students. This time-honored cloning mechanism turns AI professors into a replenishing resource like grass in a pasture – one that is prone to overexploitation.

In order to prevent the potential collapse of academic AI research, industry should be more attentive to the needs of academia. Think of academia-industry interaction as a spectrum, with all-out poaching of professors lying on one end. On the other end, some companies – including Google, Microsoft Corp., Facebook, International Business Machines Corp. and, most recently, JPMorgan Chase & Co. ⁴ – are supporting academic AI research through grants and fellowships, with no strings attached.

The most sustainable model lies between these extremes. Under this model, a professor splits her time between her home university and a company, while carrying out her usual academic responsibilities. Ideally, the company helps support the professor's research and students.

Encouragingly, several companies have been experimenting with variations of this hybrid model. Although it is fashionable to excoriate Facebook's top brass, the company's AI research division, led by New York University professor Yann LeCun, sets a positive example. In particular, Facebook has recently opened AI labs in Pittsburgh and Seattle, with the goal of tapping local professors without impeding academic research and education.

In the same vein, Google announced last month that it will open an AI lab in Princeton, New Jersey, in collaboration with Princeton University. And the Bosch Center for AI Research just opened a lab in Pittsburgh as part of a remarkable agreement with Carnegie Mellon, whereby Bosch supports AI research at the university while allowing its new chief scientist of AI research, Zico Kolter, to continue serving as a full-time faculty member.

That's a promising way to escape the oppressive logic of the tragedy of the commons. Another reason for optimism is that Lloyd's 19th-century scenario doesn't tell the whole story: AI professors typically have more autonomy than grass does. As players in the game, we can be part of the solution. In time, academics are likely to demand sustainable models of

engagement with industry. Some might even fend off the pressure to go corporate altogether, and instead strike it rich by, say, moonlighting as opinion columnists. Oh, wait.

- 1 The reader might be wondering about the missing analog of cattle. Perhaps executive recruiters?
- 2 It is convenient, but somewhat unfair, to lump the Allen Institute for AI together with the others, because its mode of operation has overall been synergistic with that of academia.
- 3 I am focusing on graduate education because, at the undergraduate level, there is only limited correlation between research prowess and teaching abilities – the best teachers are often not the most sought-after researchers – and, consequently, poaching is less disruptive at that level.
- 4 Last year, JPMorgan hired the head of the machine learning department at Carnegie Mellon, Manuela Veloso, to lead its AI research.

This column does not necessarily reflect the opinion of the editorial board or Bloomberg LP and its owners.

To contact the author of this story:

Ariel Procaccia at arielpro@cs.cmu.edu

To contact the editor responsible for this story:

Jonathan Landman at jlandman4@bloomberg.net

Ariel Procaccia is Gordon McKay Professor of Computer Science at Harvard University. His areas of expertise include artificial intelligence, theoretical computer science and algorithmic game theory.

[Read more opinion](#)