OPINION

The future of AI innovation and the role of academics in shaping it

Universities founded centuries ago are the ones poised to tackle problems far enough removed from the market demands and day-to-day profit margins that drive industry's shareholders.

By **Daniela Rus** Updated July 20, 2023, 12:28 p.m.

JÜRGEN FÄLCHLE/ADOBE

Sixty years ago, the pie-in-the-sky dream for engineers at MIT was for two people to be able to use the same computing system at the same time. That humble little initiative, which was given the unassuming title of "Project MAC" (Project Machine-Aided Cognition), later morphed into Computer Science and Artificial Intelligence Lab, or

CSAIL, which is now MIT's largest research lab, with more than 1,500 members pushing the boundaries of how computing looks and what it can accomplish.

We've evolved from an era when multiple people aspired to share a single computer platform to a time in which many of us have upward of a dozen of them. The growth has been astounding and a testament to the work being done at forward-looking universities and research labs around the world.

Far from being irrelevant, academia remains a vital hub for innovation, fostering interdisciplinary research, nurturing talent, and providing an environment conducive to long-term breakthroughs. A 2020 report found that US universities spend roughly $75 billion per year on research, roughly one-seventh of the country's entire research and development spending. Universities drive innovation by pushing the boundaries of knowledge, leading science, and training the future workforce. The National Science Foundation's Industry-University Cooperative Research Centers program has calculated that every dollar put into a university partnership by a company is leveraged 40 times.

When he was asked how to build a great city, the late Senator Daniel Patrick Moynihan of New York responded, "Create a great university and wait 200 years." This might seem incongruous with the "move fast and break things" approach of Silicon Valley, but the reality is that the supposedly slow-paced universities founded centuries ago are the ones poised to tackle problems far enough removed from the market demands and day-to-day profit margins that drive industry's shareholders. At universities, there's "no rush to deploy," allowing us to experiment in ways that lead to more unexpected future breakthroughs.

While academia's grip on tech innovation remains strong, in recent years universities have seen [a substantial decline in public funding](#) that has slowed the abilities of labs like ours to help tackle major global challenges. Take the fast-moving topic of large language models like ChatGPT. These "intelligence mimics" have put AI in everybody's pocket but still aren't well understood, and, as industry competitors race to deploy and scale, academics are the ones working to understand how these "black-box" systems work. However, as things currently stand, academia as a whole lacks industry's compute resources for answering questions about LLMs and how to ensure that they are used ethically and equitably.

Academia needs a large-scale research cloud that allows researchers to more efficiently share resources to address these sorts of hot-button issues. It would provide an integrated platform for large-scale data management, encourage collaborative studies across research organizations, and offer access to cutting-edge technologies, while ensuring cost efficiency. In the age of LLMs, the research cloud is a critical tool to propel academic research into a new era of innovation, discovery, and impact. Such unparalleled access to vast computing resources and state-of-the-art machine learning tools would empower researchers to demystify the black box, unlocking a deeper, more robust understanding of machine learning and certifiable applications.

There is much that we as computer scientists still don't understand about LLMs and the larger umbrella of generative AI structures that includes automated coding, image generation, and other energy models. AI requires massive manually labeled sets of high-quality data that need to include all of the possible types of events and failures that could happen for different applications. If the data are bad or biased, the performance of the algorithm will be, too — which is why the academic research community is working to make machine learning more trustworthy, computationally efficient, and accurate. But first we need the computational infrastructure to better be able to look under the hoods of these AI systems and understand their inner workings.

We also need to address the societal side of deploying AI systems for the greater good. The spread of AI has the potential to make our lives easier by easing the burden of many of our dull, dirty and dangerous tasks, but some of the roles that it can play will displace work done by humans. We need to anticipate and respond to the economic inequality this could create. The lack of interpretability and dependence also leads to significant issues around trust and privacy, requiring a more robust ethical and legal framework for AI. These problems are ones that we know are coming, which means that academic researchers can be proactive in setting out to find solutions.

If we don't empower universities with a large-scale research cloud to enable us to better study and understand machine learning, we run the risk of creating a future where we cannot fully control the technological innovations that we have invented.

Daniela Rus is director of CSAIL and a professor of electrical engineering and computer science at MIT.

Show comments