Moshe Y. Vardi

DOI:10.1145/3352489

The Long Game of Research

THE INSTITUTE FOR THE FUTURE (ITFF) in Palo Alto, CA, is a U.S.-based think tank. It was established in 1968 as a spin-off from the RAND Corporation to help organizations plan for the long-term future. Roy Amara, who passed away in 2007, was ITFF's president from 1971 until 1990. Amara is best known for coining Amara's Law on the effect of technology: "We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run." This law is best illustrated by the Gartner Hype Cycle,^a characterized by the "peak of inflated expectations," followed by the "trough of disillusionment," then the "slope of enlightenment," and, finally, the "plateau of productivity."

I was reminded of Amara's Law when I heard that the 2018 Turing Award was awarded to Yoshua Bengio, Geoffrey Hinton, and Yann LeCun for "conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing." This decision was hardly surprising. After all, it is difficult to think of any other computing technology that has such a dramatic appearance and impact over the past decade. Quoting the Turing Award announcement: "In recent years, deep-learning methods have been responsible for astonishing breakthroughs in computer vision, speech recognition, natural language processing, and robotics—among other applications."

But it is worthwhile to reflect on the long history of neural nets in order to put this contribution in its proper historical context. In 1943, Warren McCulloch, a neurophysiologist, and a young mathematician, Walter Pitts, wrote a paper on how brain neurons might work. They modeled a simple

neural network with electrical circuits. Frank Rosenblatt, a neurobiologist of Cornell, invented the *Perceptron*, a single-layer neural net, in 1958. *The New York Times* reported the perceptron to be "the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence." Unfortunately, the perceptron is quite limited and was proven as such in Marvin Minsky and Seymour Papert's 1969 book, *Perceptrons*. The peak of hype was then followed by the trough of disillusionment. This so-called "First AI Winter" manifested, among other things, in the declining research funding for artificial intelligence, and lasted until the early 1980s.

In 1982, John Hopfield of Caltech presented a paper with a focus not on modeling brains but on creating useful devices. With mathematical clarity, he showed how such networks could work and what they could do. Around the same time, a *U.S.-Japan Joint Conference on Cooperative/Competitive Neural Networks* was held in Kyoto, Japan. Japan subsequently announced its Fifth Generation effort. U.S. periodicals picked up that story, generating a worry that the U.S. could be left behind. Soon funding was flowing once again. The Annual Conference on Neural Information Processing Systems was launched in 1987. Yet the new peak of hype was again followed by a trough of disillusionment. Quoting again the Turing Award announcement: "By the early 2000s, LeCun, Hinton, and Bengio were among a small group who remained committed to this approach." In fact, their efforts to rekindle the AI community's interest in neural networks were initially met with skepticism. This disillusionment led to the "Second AI Winter," which lasted well into the 1990s.

It was only at the start of this decade that the combination of improved algorithms, improved hardware (GPUs), and very large datasets (ImageNet has more than 14 million labeled images) led to an impressive breakthrough, and it became obvious that deep (many layered) neural nets had significant advantages for machine vision, in terms of efficiency and speed. The ideas of Hinton and his colleagues resulted in major technological advances, and their methodology is now the dominant paradigm in the field, leading to being awarded the 2018 Turing Award.

The moral of this tale is that research is a long game; patience and endurance are necessary components. Yet I remember a research-evaluation meeting in an industrial-research lab in the early 1990s in which someone's seminal work on data mining was not being appreciated, because "he has been doing it for two years now and it is not clear that it is going anywhere." I share the concerns of Abraham Flexner, founder of the Institute for Advanced Study in Princeton; in *The Usefulness of Useless Knowledge*,^b published 1939, Flexner explores the dangerous tendency to forgo pure curiosity in favor of alleged pragmatism.

There is no single formula for successful research. Sometimes it makes sense to focus short-term on an immediate problem, but, quite often, dramatic breakthroughs are obtained by viewing research as a long game.

Follow me on Facebook and Twitter. ☐

^b <https://library.ias.edu/files/UsefulnessHarpers.pdf>

Moshe Y. Vardi (vardi@cs.rice.edu) is the Karen Ostrom George Distinguished Service Professor in Computational Engineering and Director of the Ken Kennedy Institute for Information Technology at Rice University, Houston, TX, USA. He is the former Editor-in-Chief of *Communications*.

Copyright held by author.

^a http://en.wikipedia.org/wiki/Hype_cycle