THREE MAJOR TRENDS IN AI RESEARCH TECHNIQUES

AI research has long been characterized by the sudden rise and fall of different techniques, and every decade has seen a heated competition between different ideas.

There is a study based on the abstracts of all 16,625 papers available in the "artificial intelligence" section ofone of the largest open-source databases of scientific papers, known as the arXiv.

The results show three major trends: **a shift toward machine learning during the late 1990s and early 2000s, a rise in the popularity of neural networks beginning in the early 2010s, and the growth in** reinforcement learning **in the past few years.** Every decade has essentially seen the reign of a different technique: neural networks in the late 1950s and 1960s, various symbolic approaches in the 1970s, knowledge-based systems in the 1980s, Bayesian networks in the 1990s, support vector machines in the 2000s, and neural networks again in the 2010s.

The first shift

The first shift happened in the early 2000s, when knowledge-based systems phased out in favor of machine learning. The former computer programs are based on the idea that you can use rules to encode all human knowledge. The latter are the parent category of algorithms that includes deep learning.

In the 1980s, knowledge-based systems amassed a popular following thanks to the excitement surrounding ambitious projects that were attempting to re-create common sense within machines. But as those projects unfolded, researchers hit a major problem: there were simply too many rules that needed to be encoded for a system to do anything useful. This jacked up costs and significantly slowed ongoing efforts.

Figure 1

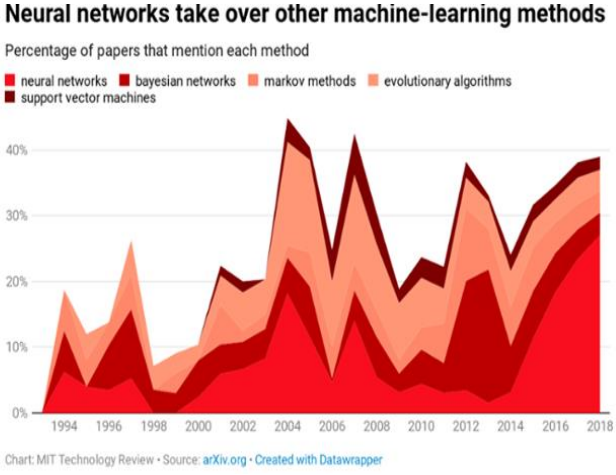

Machine learning became an answer to that problem. Instead of requiring people to manually encode hundreds of thousands of rules, this approach programs machines to extract those rules automatically from a pile of data. Just like that, the field abandoned knowledge-based systems and turned to refining machine learning.

The second shift

The second shift—a convergence of machine learning to deep learning—happened in the early 2010s. But for the first decade after the rise in popularity of machine learning, researchers tested a variety of methods for finding patterns in data.

Competition stayed steady through the 2000s before a pivotal breakthrough in 2012 led to another sea change. During the annual ImageNet competition, intended to spur progress in computer vision, a researcher named Geoffrey Hinton, along with his colleagues at the University of Toronto, achieved the best accuracy in image recognition by an astonishing margin of more than 10 percentage points.

The technique he used, machine learning through deep neural networks, aka: deep learning, sparked a wave of new research within the vision community.

The third shift

The third shift happened around the same time as the upsurge in deep learning. As well as the different techniques in machine learning, there are three different types: supervised, unsupervised, and reinforcement learning. Supervised learning, which involves feeding a machine labeled data, is the most commonly used and also has the most practical applications by far. In **the last few years, however, reinforcement learning, which mimics the process of training animals through punishments and rewards,** has seen a rapid uptick of mentions in paper abstracts.

The idea isn't new, but for many decades it didn't really work. But, just as with deep learning, one pivotal moment suddenly placed it on the map. That moment came in October 2015, when DeepMind's AlphaGo, trained with reinforcement learning, defeated the world champion in the ancient game of Go. The effect on the research community was immediate.

Compiled by SCM (2019) from

Karen Hao, January 25, 2019, MIT Technology Review