

Fact vs. futuristic fiction:

The realities of MACHINE LEARNING

BY MADISON MOORE

In the late 19th century, the Industrial Revolution introduced complicated farm machinery that changed the way farmers planted, cultivated and harvested their crops. These machines meant fewer farm hands were needed, but it created jobs for people to build and repair farm equipment. The improved benefits of utilizing farm machinery impacted the quality of life, it produced food faster, and it created jobs and a new life for farmers.

Machine learning is today's Industrial Revolution. TV shows and movies currently portray machine learning as this creepy, self-aware, futuristic technology that takes over humans' jobs, but these examples do not properly show the real advancements of these cognitive systems. The real improvements are like that of farm machinery; these systems will work together with humans side by side, creating experiences that have nev-

er been attainable before.

Machine learning is not "new," but in order to understand how far the industry has come, it helps to take a step back and see where the technology and research began.

It was in the 1950s when machine learning research was conducted using simple algorithms. Around this time, Alan Turing proposed his "learning machines," and he created the Turing Test to determine if computers have real intelligence. Following this achievement in 1957, psychologist Frank Rosenblatt invented the perceptron, or an algorithm for supervised learning of binary classifiers.

In 1980, research scientist Kunihiko Fukushima published work on a type of neural network, which later inspired research on convolutional neural networks, or neural networks that can be applied to visual recognition tasks. The

rest of the 1980s was filled with other machine learning discoveries, like research on recurrent neural networks, backpropagation, and reinforcement learning.

In 1992, principal researcher at IBM's TJ Watson Research Center Gerald Tesauro developed a computer backgammon program that used an artificial neural network. It became a rival to human backgammon players. Then, IBM's Deep Blue beat the world champion at chess in 1997, and most recently, Google's AlphaGo became the first Computer Go program to beat a professional Go player.

What is machine learning?

The previously highlighted milestones led to successful scenarios of machine learning doing things like recommending items to consumers on platforms like Amazon, Netflix, and Spotify. That

technology has led to such things as spam filters, speech recognition systems, retargeted Internet advertising, and machine translations.

Today, machine learning is easier and developers or researchers can tap into power that didn't exist before. We can hardly go a few months without seeing a new deep learning architecture or machine learning framework that claims it supersedes all previous standards, said CEO of Bonsai, Mark Hammond. Bonsai centralizes the programming and management of AI models into one platform.

Machine learning is a word that gets tossed around, but Steve Abrams, vice president of developer advocacy for IBM Watson, said machine learning is simply a set of algorithms whose behavior is determined by experiences.

Here's an example: Imagine you want to build a computer vision system that can recognize trees, he said. With machine learning, you can build a system that can be fed examples of pictures of trees and non-trees. Over time, the computer determines what features and patterns depict a tree, said Abrams, so eventually, the system can recognize what is a tree and what isn't.

"Machine learning is this set of capa-

model, which represents the data-generating process. Essentially, machine learning takes raw data and it tries to predict something from that data, he said.

This raw data comes in handy for companies that want to predict if a customer is going to leave and go to a rival, for example. According to Gualtieri, the company would use machine learning on historical data to determine a predictive model, which would give them valuable business insights, he said.

"If you know within 89% probability if that customer is likely to leave you, then you have time to do something about it," said Gualtieri. "If you think about the recommendations on Netflix, those are done with machine learning and that's a predictive model as well; it's predicting what TV or movie you are most likely to enjoy watching. Machine learning is all about creating that predictive model."

The reality of machine learning

There have been examples of machine learning being able to do things like flip burgers and write novels, but the realities of machine learning extend beyond just "cool" capabilities.

Think of machine learning as the "latest new old thing," said Phil Tee,

where data is stored and where it lives, and who is going to be the company to deploy it, he said.

Right now, the majority of data for enterprises lives on-premises in an Oracle, IBM or SAP ecosystem, and this has created an almost war-like effort of who will win the "infrastructure battle," said Liu. Because of all the data that exists, companies have to consider if it is cost-efficient to buy all the gear and host the data on premises, or if it makes more sense to rent the equipment or have a cloud-based model.

Besides the role data plays today, the other notable change is the power of the platform, said Tee. He also sees advancements in voice recognition, image recognition, statistical translation, and semantic indexing of knowledge. Tee believes that we are headed into a digital assistant future, and the next 10 to 15 years will consist of advancements in self-driving vehicles and other digital assistants in applications, homes, and other devices.

The biggest advancement, according to Abrams, is that machine learning is not just a theory; it's a set of mature technologies that can actually be applied to the real world, like in healthcare, applications, automobiles, and hospitality sce-

Today, machine learning is easier and developers or researchers can tap into power that didn't exist before.



bilities; it can learn models that can predict the future, it can learn models that are going to interpret sensor data, and each of those models are shaped by the experiences they've been given," said Abrams.

Another way to translate machine learning and understand its purpose is to think about predictive models or probabilistic machine models, said Mike Gualtieri, vice president and principal analyst at Forrester Research. Predictive modeling uses data mining and probabilistic models to formulate a statistical

CEO of Moogsoft, an IT operations analytics company. It's not news that businesses are using machine learning, but the important thing to note is that now, the technology is being used at a bigger scale.

The big difference of machine learning today is the role that data plays in its algorithms. It's true, according to Kerry Liu, CEO of Rubikloud Technologies, a machine intelligence platform for enterprise retailers, that the industry can't talk about machine learning without mentioning the underlying shift in

narios. Years ago, these technologies were applied in much more narrow domains, and now it's readily available for developers, researchers, and other to consume, he said.

"Go back in time, and you pretty much needed a PhD in machine learning and you needed a PhD in data science to make progress with these technologies," said Abrams. "Because of the way we've been able to encapsulate [the technology] and make it available as APIs that are consumable by mere mor-

continued on page 37 ►

◀ continued from page 35

tal developers, if you will, we are really kind of democratizing the access to the capabilities.”

This means developers without extensive background knowledge in data science can actually crowdsource the training of a vision model and produce a vision model to “do amazing things,” said Abrams. For instance, one of IBM’s partners, OmniEarth, was able to build a model that can recognize landscape features and building features, and use it to predict water consumption. This was done using aerial photography and IBM machine learning technology, he said.

These examples extend into the enterprise world, too, according to Rubikloud’s Liu, but enterprises are not the companies buying the algorithms

that are available today, they are not buying the machine learning libraries, nor are they implementing them at a practical level, he said.

Enterprises are waiting for the software to come out, whether it be from Google, Oracle, Microsoft, or startups, and they want to buy machine learning software that solves their day-to-day problems. The also want to make sure it connects to their ROI, said Liu.

“If I buy a piece of software, which now has the machine learning technology embedded in it, what part of my P&L is going to be driven? Am I going to save money by not buying that other piece of software? Am I going to make more money because my customers are more optimized?” said Liu. “These companies are trying to connect the machine learning hype with the reali-

ties of their business, and until you can prove that connection to them, you are not going to sell it to them.”

Getting machine-learning savvy

According to Hammond, the shortage of talent for data scientists and developers who can create high-level machine learning algorithms is very real. For companies that are competing against the big names like Google and Amazon, who are all fighting to get the best data scientists, it becomes more of an issue of availability of the resources that have the talent to do the actual work, he said.

Machine learning used to be an obscure elective, but now, many developers are graduating with enough knowledge about machine learning to get out there and start doing something

continued on page 38 ▶



Machine learning isn’t the only term getting all the buzz. Deep learning, or a class of machine learning algorithms, is showing great promise, primarily because it’s getting results.

Deep learning historically was largely inaccessible because it had such high demand on computational resource and data, but with the progression of technology, storage costs have come down and the computation has gone up, said CEO of Bonsai, Mark Hammond. It demands a lot of resources, but there are a few organizations focusing on making this technology and research accessible and easy to use (think: OpenAI or Google’s DeepMind).

According to vice president and principal analyst at Forrester Research, Mike Gualtieri, machine learning is trying to find a predictive model, whereas deep learning is based upon a hierarchical network, roughly fashioned after the human brain, he said. Since it’s modeled after the brain, people argue that deep learning actually lets systems “learn,” said Gualtieri.

“Does it learn the way we learn? Roughly,” said Gualtieri. “The reality is we are not sure. It’s a very narrow scope; you can train a machine, but it’s always about predicting one thing. Our brain and your brain, we can do much more than just predict the next move someone likes. It’s not generalized learning.”

Deep learning is a subfield of machine learning based on neural networks, which was first written about in 1943 by neurophysiologist Warren McCulloch and mathematician Walter Pitts. In order to describe how neurons in the brain might work, they modeled simple neural networks using electrical circuits. According to Gualtieri, deep learning neural networks were hard to train, and around the year 2012, there was a research break-

through that made it very practical to do deep learning. The reason why deep learning is just now getting all the attention is because it’s making new state-of-the-art results.

“The fact that these technologies are easier to apply, easier to scale, easier to get good results, means there is just a huge range of applicability,” said chief application architect at MapR, Ted Dunning. “Maybe even in three years it will be a standard part of a developer’s toolkit to build simple [deep learning] models and integrate them into their systems.”

Machine learning, and specifically deep learning, has allowed researchers to do a number of things substantially better than before, said Ash Munshi, CEO of Pepperdata.

For instance, using deep learning, we can understand images, speech, translate language, play games, and build self driving cars — all of these advancements were not possible just a few years ago, said Munshi.

Deep learning can be used to review an image and generate a textual description of the image. It can be used to translate from one language to another accurately, and it can even speak text naturally without “any of that robotic tone that we are all so accustomed to hearing,” added Munshi.

“Machines with deep learning capabilities will likely replace call center personnel, personal assistants, physician assistants and a whole host of other such tasks and jobs,” said Munshi. “They will simply be better and more accurate than humans in these regards, and since they will be natural to use, they will become omnipresent. Deep learning will be like electricity, it will be everywhere.” ■

—Madison Moore

◀ continued from page 37

with existing tools, said Gualtieri.

However, machine learning and deep learning are not tied into most average computer science curriculums, said chief application architect at MapR, Ted Dunning. Deep learning is not the skill, he said, data orientation is the skill.

“[Developers need] to see where data might come from, see when it is corrupted or valueless, have the ability to build simple visualizations to see how data is connected to things, and be able to prove simple hypotheses from data,” said Dunning.

There are also pretrained models in the cloud, said Gualtieri, where developers can tap into available data from companies like Microsoft or IBM, and they can access it as API calls on the Internet. So if a developer wants to do an image analysis, they can call the AWS Rekognition service, for example. The developer feeds the service an image, it has a machine learning model, and it displays the objects with keywords from the actual photo, he said.

Anyone who wants to get started with machine learning or deep learning needs to make sure they have a rich enough and large enough data set, said Ash Munshi, CEO of Pepperdata, a Big Data performance company. Those that do not have a large data set will not get any worthwhile results, he said.

“It is also extremely important that you understand how to train and test your models. In most cases, simple techniques are all that are needed, [so] look at examples and find the best matches,” said Munshi.

Developers can also incorporate some machine learning capabilities by using APIs, and the more ambitious developers can use open-source tools to create the models themselves, said Gualtieri. He said there are a lot of mathematical tests that developers need to apply to a model in order to make sure it works, and this is getting easier with tools since they have those built-in capabilities.

“There will be new skills [developers need] but they are not that hard to pick up if you have basic science backgrounds; the idea that you measure the

Getting started with machine and deep learning

There are plenty of open-source frameworks for deep learning and machine learning, but choosing one to implement depends on the goals and problem at hand. Here are some popular deep learning and machine learning frameworks and projects found on the Internet and recommended by some of our machine learning experts:

■ **Keras:** A high-level neural networks API written in Python. It can run on top of TensorFlow or Theano. Use Keras for a deep learning library that supports fast prototyping, convolutional networks and recurrent networks, and if you want to work in Python code.

■ **TensorFlow:** This popular open-source deep learning framework leverages Google’s infrastructure for scalable training. It provides rich higher level tools for language, image and video understanding.

■ **Theano:** A Python library that lets developers define and evaluate mathematical expressions involving multi-

dimensional arrays. It uses GPUs and performs efficient symbolic differentiation.

■ **Mxnet:** An open-source framework built by Amazon. Mxnet produces models that are compact, require less memory and CPU for inference.

■ **Caffe:** A deep learning framework made with expression, speed and modularity in mind. Speed is a big feature in Caffe; it can process over 60 million images per day with a single NVIDIA K40 GPU.

■ **Sonnet:** This is a library built on top of TensorFlow (TF) for building complex neural networks. Google’s DeepMind is developing the codebase for building neural network modules with TF. Models written in Sonnet can be freely mixed with raw TF code and other high-level libraries.

■ **Scikit Learn:** This is a toolkit for doing classical machine learning, since not all problems need the power of deep learning. Scikit provides a rich set of tools for data mining and data analysis.

real world, you get data and you draw conclusions from that data,” Dunning said.

When machines rule the world

Fear not machine learning skeptics, the world will not be filled with robots who plan to take over the world. But, we would be deluding ourselves to say machine learning will not impact the workforce, said Bonsai’s Hammond.

To some extent, the hype that machine learning will take over jobs is “silly,” said Forrester’s Gualtieri. Some jobs of course are very narrow and manual, and these are positions that would benefit from machine learning technology. But in cases where careers require a great deal of creativity, there’s no way to build that into a machine, he said.

“It’s augmenting the abilities of the humans that are doing the work, and allowing them to do more and do it better as opposed to replicating them,” said Hammond. “That is almost always the path the technological progress follows.”

The bottom line is, machine learning will make humans more productive,

especially with the rise of “personal assistants,” or systems that take on some of the manual, tedious work that humans get bogged down with.

Physicians might spend half their day on paperwork or manual tasks, but imagine a machine learning assistant that could do all of this work and explain the prognosis in 17 minutes, said Abrams. This would allow the doctor to spend more time doing what they need to do, like meeting with patients.

In the future, one of the things the industry will see is the focus on the emotional connection that people develop with machine learning assistants, said Abrams. He’s not referring to the chilling sci-fi scenarios depicted in movies; he means building systems that don’t just understand the meaning of word, but the context of the word.

“When I look down the road, I think that systems that understand people are going to be increasingly important,” said Abrams. “AI systems are here to work with us; not take over us.” ■

Read this story on
sdtimes.com

