AI researcher offers insight on promise, pitfalls of machine learning

29 July 2019, by Victor Chen

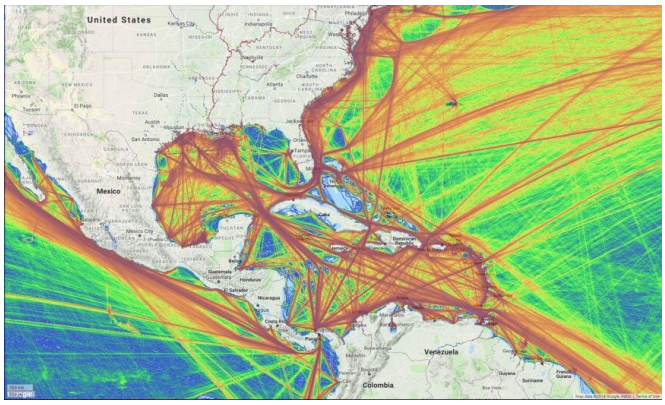

MarineTraffic's Density Map format showing vessel trajectories from billions of data points from 2017. The 'cool' colored lines signify that a route has not been taken often, the 'warm' colored lines signify where routes are often utilized. The result is a global dataset of ship tracking density. Credit: MarineTraffic

These days, the latest developments in artificial intelligence (AI) research always get plenty of attention, but an AI researcher at the U.S. Naval Research Laboratory believes one AI technique might be getting a little too much.

Ranjeev Mittu heads NRL's Information Management and Decision Architectures Branch and has been working in the AI field for more than two decades.

"I think people have focused on an area of machine learning—deep learning (aka deep networks)—and less so on the variety of other artificial intelligence techniques," Mittu said. "The biggest limitation of deep networks is that a complete understanding of how these networks arrive at a solution is still far from reality."

Deep learning is a machine learning technique that can be used to recognize patterns, such as identifying a collection of pixels as an image of a dog. The technique involves layering neurons together, with each layer devoted to learning a different level of abstraction.

In the dog image example, the lower layers of the neural network learn primitive details like pixel values. The next set of layers attempt to learn edges; higher layers learn a combination of edges as a nose. With enough layers, these networks are able to recognize images with nearly human-like performance.

But the systems can be fooled easily just by changing a small number of pixels, according to Mittu.

"You can have adversarial 'attacks' where once you've created a model to recognize dogs by showing it millions of pictures of dogs," he said. "…making changes to a small number of pixels, the network may misclassify it as a rabbit, for example."

The biggest flaw in this machine learning technique, according to Mittu, is that there is a large degree of art to building these kinds of networks, which means there are very few scientific methods to help understand when they will fail.

The solution?

"There are numerous AI techniques of which machine learning is a subset," he said. "While deep learning has been highly successful, it is also currently limited because there is little visibility into its decision rationale. Until we truly reach a point where this technique becomes fully "explainable," it cannot inform humans or other automation as to how it arrived at a solution, or why it failed. We have to realize that deep networks are just one tool in the AI tool box."

And, humans have to stay in the loop.

"Imagine you have an automated threat detection system on the bridge of your ship, and it picks up a small object on the horizon," he said. "The deep network classification may indicate that it is a fast attack craft coming at you, but you know that a very small set of uncertain pixels can mislead the algorithm. Do you believe it?

"A human will have to examine it further. There may always need to be a human in the loop for high risk situations. There could be a high degree of uncertainty and the challenge is to increase the classification accuracy while keeping the false alarm rate low—it is sometimes very difficult to strike the perfect balance. "

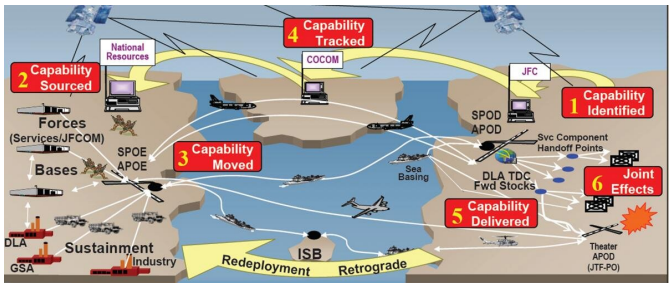

Integrated Data Environment and Global Transportation Network Convergence (IGC). Credit: U.S. Transportation Command/Defense Logistics Agency

The problem with machine learning

When it comes to machine learning, the key factor, simply put, is data.

Consider one of Mittu's previous projects: an analysis of commercial shipping vessel movements around the world. The project's goal was to use machine learning to discern patterns in vessel traffic to identify ships involved in illicit activities. It proved a difficult problem to model and understand using machine learning, Mittu said.

"We cannot have a global model because the behaviors will differ for vessel classes, owners, etc." he explained. "It is even different seasonally, because of sea state and weather patterns."

But the bigger problem, Mittu found, was the possibility of mistakenly using poor quality data.

"Ships transmit their location and other information, just like aircraft. But what they transmit can be spoofed," Mittu said. "You don't know if it is good or bad information. It is like changing those small numbers of pixels on the dog image that causes the system to fail."

Missing data is another issue. Imagine a case in which you must move large numbers of people and materials on a regular basis to sustain military operations, and you're relying on incomplete data to predict how you might act more efficiently.

"The difficulty comes when you start to train machine learning algorithms on data that is of poor quality," Mittu said. "Machine learning becomes unreliable at some point, and operators will not trust the outcomes of the algorithms."

Current work in AI

Today Mittu's team continues to pursue AI innovations in multiple areas of the field. They advocate an interdisciplinary approach to employing AI systems to solve complex problems.

"There are many ways to improve predictive capabilities, but probably the best–of-breed will take a holistic approach and employ multiple AI techniques and strategically integrate the human decision-maker," he said.

"Aggregating various techniques (similar to 'boosting'), which may 'weight' algorithms differently, could provide a better answer, or learning combined with reasoning, etc. By employing combinations of AI techniques, the resulting system may be more robust to poor data quality."

One area Mittu is excited about is recommender systems. According to him, most people are already familiar with these systems, which are used in search engines and entertainment applications such as Netflix. He's excited about the potential military applications.

"Think of a military command and control system, where users need good information to make good decisions," he said. "By looking at what the user is doing in the system within some context, can we anticipate what the user might do next and infer the data they might need."

While the field of AI offers almost limitless potential for innovative solutions to today's problems, Mittu said, researchers obviously have many years of work ahead of them.

"We need to determine the right techniques, their limitations, and the data that is needed in order to get reliable answers in order for the users to trust the resulting system," he said. "The field of AI has a long way to go in taking a holistic approach by strategically integrating the decision-maker in order to improve the performance of the human and machine system."

Provided by Naval Research Laboratory