

The uncertain unicycle that taught itself and how it's helping AI make good decisions

February 15 2018, by Louise Walsh



Credit: The District

Cambridge researchers are pioneering a form of machine learning that starts with only a little prior knowledge and continually learns from the world around it.

In the centre of the screen is a tiny unicycle. The animation starts, the unicycle lurches forward and falls. This is trial #1. It's now trial #11 and there's a change – an almost imperceptible delay in the fall, perhaps an

attempt to right itself before the inevitable crash. "It's learning from experience," nods Professor Carl Edward Rasmussen.

After a minute, the unicycle is gently rocking back and forth as it circles on the spot. It's figured out how this extremely unstable system works and has mastered its goal. "The unicycle starts with knowing nothing about what's going on – it's only been told that its goal is to stay in the centre in an upright fashion. As it starts falling forwards and backwards, it starts to learn," explains Rasmussen, who leads the Computational and Biological Learning Lab in the Department of Engineering. "We had a real unicycle robot but it was actually quite dangerous – it was strong – and so now we use data from the real one to run simulations, and we have a mini version."

Rasmussen uses the self-taught unicycle to demonstrate how a machine can start with very little data and learn dynamically, improving its knowledge every time it receives new information from its environment. The consequences of adjusting its motorised momentum and balance help the unicycle to learn which moves were important in helping it to stay upright in the centre.

"This is just like a human would learn," explains Professor Zoubin Ghahramani, who leads the Machine Learning Group in the Department of Engineering. "We don't start knowing everything. We learn things incrementally, from only a few examples, and we know when we are not yet confident in our understanding."

Ghahramani's team is pioneering a branch of AI called continual [machine learning](#). He explains that many of the current forms of machine learning are based on neural networks and deep learning models that use complex algorithms to find patterns in vast datasets. Common applications include translating phrases into different languages, recognising people and objects in images, and detecting unusual

spending on credit cards.

"These systems need to be trained on millions of labelled examples, which takes time and a lot of computer memory," he explains. "And they have flaws. When you test them outside of the data they were trained on they tend to perform poorly. Driverless cars, for instance, may be trained on a huge dataset of images but they might not be able to generalise to foggy conditions.

"Worse than that, the current deep learning systems can sometimes give us confidently wrong answers, and provide limited insight into why they have come to particular decisions. This is what bothers me. It's okay to be wrong but it's not okay to be confidently wrong."

The key is how you deal with uncertainty – the uncertainty of messy and missing data, and the uncertainty of predicting what might happen next. "Uncertainty is not a good thing – it's something you fight, but you can't fight it by ignoring it," says Rasmussen. "We are interested in representing the uncertainty."

It turns out that there's a mathematical theory that tells you what to do. It was first described by 18th-century English statistician Thomas Bayes. Ghahramani's group was one of the earliest adopters in AI of Bayesian probability theory, which describes how the probability of an event occurring (such as staying upright in the centre) is updated as more evidence (such as the decision the unicycle last took before falling over) becomes available.

Dr. Richard Turner explains how Bayes' rule handles continual learning: "the system takes its prior knowledge, weights it by how accurate it thinks that knowledge is, then combines it with new evidence that is also weighted by its accuracy.

"This is much more data-efficient than the way a standard neural network works," he adds. "New information can cause a neural network to forget everything it learned previously – called catastrophic forgetting – meaning it needs to look at all of its labelled examples all over again, like relearning the rules and glossary of a language every time you learn a new word.

"Our system doesn't need to revisit all the data it's seen before – just like humans don't remember all past experiences; instead we learn a summary and we update it as things go on." Ghahramani adds: "The great thing about Bayesian machine learning is the system makes decisions based on evidence – it's sometimes thought of as 'automating the scientific method' – and because it's based on probability, it can tell us when it's outside its comfort zone."

Ghahramani is also Chief Scientist at Uber. He sees a future where [machines](#) are continually learning not just individually but as part of a group. "Whether it's companies like Uber optimising supply and demand, or autonomous vehicles alerting each other to what's ahead on the road, or robots working together to lift a heavy load – cooperation, and sometimes competition, in AI will help solve problems across a huge range of industries."

One of the really exciting frontiers is being able to model probable outcomes in the future, as Turner describes. "The role of uncertainty becomes very clear when we start to talk about forecasting future problems such as climate change."

Turner is working with climate scientists Dr. Emily Shuckburgh and Dr. Scott Hosking at the British Antarctic Survey to ask whether machine learning techniques can improve understanding of climate change risks in the future.

"We need to quantify the future risk and impacts of extreme weather at a local scale to inform policy responses to climate change," explains Shuckburgh. "The traditional computer simulations of the climate give us a good understanding of the average climate conditions. What we are aiming to do with this work is to combine that knowledge with observational data from satellites and other sources to get a better handle on, for example, the risk of low-probability but high-impact weather events."

"It's actually a fascinating machine learning challenge," says Turner, who is helping to identify which area of climate modelling is most amenable to using Bayesian probability. "The data are extremely complex, and sometimes missing and unlabelled. The uncertainties are rife." One significant element of uncertainty is the fact that the predictions are based on our future reduction of emissions, the extent of which is as yet unknown.

"An interesting part of this for policy makers, aside from the forecasting value, is that you can imagine having a machine that continually learns from the consequences of mitigation strategies such as reducing emissions – or the lack of them – and adjusts its predictions accordingly," adds Turner.

What he is describing is a machine that – like the unicycle – feeds on uncertainty, learns continuously from the real world, and assesses and then reassesses all possible outcomes. When it comes to climate, however, it's also a machine of all possible futures.

Provided by University of Cambridge

Citation: The uncertain unicycle that taught itself and how it's helping AI make good decisions (2018, February 15) retrieved 26 April 2024 from

<https://techxplore.com/news/2018-02-uncertain-unicycle-taught-ai-good.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.