tsg insights

November 2021

Testing of artificial intelligence (AI) and machine learning (ML) –reinforcement learning

This whitepaper is the fifth in a series which act as companion pieces:

- An overview of artificial intelligence (AI) and machine learning (ML).
- Introduction to the testing of AI and ML.
- Testing of AI (artificial intelligence) and ML (machine learning) – supervised learning.
- Testing of AI (artificial intelligence) and ML (machine learning) – unsupervised learning.
- Testing of AI (artificial intelligence) and ML (machine learning) – reinforcement learning.
- AI and machine learning – algorithmic bias – the cruel mirror AI and ML reflects back at us.

In this paper, we look at the approaches to and challenges regarding reinforcement learning.

What is reinforcement learning?

This machine learning method is based on rewards and/or punishments. Optimal behaviour or action by a machine is reinforced by a positive reward. Undesired ones are punished – not just by lack of rewards but by terminating a procedure (for example, in game playing that would mean ending the game). In this way reinforcement learning algorithms determine the best behaviour based upon feedback they receive from their environment.

It's been compared to a small child learning how to walk. If a toddler takes a big step and falls over, then next time they may take a smaller step until they become more adept and confident.

It's not just another way of saying supervised learning, which is based on using historical examples to make sense of an environment. For example, driving in traffic one day could radically differ from what it was the day before. You therefore need more than just remembered observations, and this is an example where reinforcement learning is different from – and here superior to – supervised learning.

There's also what's known as deep reinforcement learning. Rather than learning by trial and error to maximise outcome, this is about learning from existing knowledge and applying it to a new dataset.

Practical applications of reinforcement learning

Autonomous cars – taking the above driving example further, reinforcement learning allows for highly complex and varying scenarios. For example, the weighing up of safety, comfort, the law, emissions, time and speed. It's simply not possible to predict and thus build-in every single possible scenario and how it might unfold in real life. However, by giving different rewards and penalties, the system can learn how to reach desired goals in circumstances which cannot be totally foreseen. As well as cars, most autonomous lorries, ships and drones are being developed with the aid of reinforcement learning.

Prosthetics – a "Learning to run" project trained a virtual runner – an advanced musculoskeletal model – as the first step in developing a new generation of prosthetic legs. People differ in their walking patterns by often minute amounts, and reinforcement learning can be used to make movement more effective and natural in these circumstances.

Practical applications of reinforcement learning

Game playing – one of the first and best-known applications of reinforcement learning was when the machine learning algorithm AlphaGo played Go with one of the world's best human players and won. Reinforcement learning is now used in all kinds of games, both in their development and actual playing.

Medicine – reinforcement learning has many applications in healthcare, from clinical trials through to creating optimal treatments and drug combinations for health conditions. That includes the discovery and generation of dynamic treatment regimes (DTRs) for chronic diseases. It can also be used to quantifying the effects of delaying treatment.

Industry – there are numerous industrial applications for reinforcement learning. These include building intelligence into energy, manufacturing, automotive and supply chains, as well as the optimisation of predictive maintenance. One well-known HVAC example is enabling Google to optimise cooling requirements in their data centres – this has enabled them to reduce energy spending by 40%.

Financial trading – reinforcement learning models can be employed to determine whether to buy, sell or hold stocks based on their market price and other factors. Previously, analysts had to make decisions themselves – not easy given the multiplicity of stocks and the constant fluctuation of market prices.

Online recommendations – while unsupervised learning is good for categorising and prioritising news articles, reinforcement learning goes a step further by tracking reader behaviour to make more accurate recommendations. A wider application is where targeted advertising is delivered to individuals based on their previous purchases or browsing habits.

Leisure – the artificial intelligence research laboratory OpenAI was set up in 2015 to develop "friendly AI" that benefits society. Its research focuses on reinforcement learning, with applications including motor skills, music and gaming. It also supports individuals in creative activities such as writing and composing.

Is reinforcement learning the future of all machine learning?

Reinforcement learning offers immense potential increases in efficiency across manufacturing and supply chains, as well as environmental benefits such as decreasing carbon emissions. Deep reinforcement learning in particular is producing results that consistently outperform supervised and unsupervised learning systems.

The response to the COVID pandemic has demonstrated what can be done in a short space of time, from mapping to projections, as well as the development and production of vaccines. All the same, these advances have been achieved by a combination of unsupervised, supervised and reinforcement learning approaches.

Other machine learning approaches will continue to be used, and to improve. Reduced training time is one key area where reinforcement learning will continue to defer to supervised learning, while the need for less data input ensures that unsupervised learning will also be preferred for certain applications. In addition, many applications simply don't require the advanced levels of sophistication associated with reinforcement learning.

All the same, to make a machine creative, in the sense of seeking new ways to perform its tasks and meet goals, reinforcement learning is essentially the way ahead. All major tech companies have invested heavily in reinforcement learning research, while new companies are springing up to offer innovative products and to broaden potential applications. So while reinforcement learning isn't the whole future of machine learning, it's set to be the major part of it.