Reinforcement Learning in Industrial Applications

Eric Bergvall and Niklas Kotarsky

September 2020

1 A Computer Running the Control Room?

In the last couple of years, we have seen computer algorithms beat human players in computer games and board games. These computer algorithms are called reinforcement learning algorithms. The most famous example is Deepmind's algorithm Alpha Go which beat Lee Sedol from South Korea in the ancient game of Go which has more possible board positions than atoms in the universe. [1] [2]

To understand where challenges arise one must understand the *Reinforcement* in reinforcement learning. What it means is that the algorithm learns by playing many, many games and thus *reinforcing* positive behaviour. This causes a challenge for scientists when they want to apply these superhuman achievements of reinforcement learning into industrial settings. How should one let an algorithm play a large number of games in an industrial process? The algorithms can be likened to a child exploring the world. A *very* stubborn and curious child. As in the analogy with a child, it is very important for the algorithms to explore bad and even catastrophic behaviour and learn from their failures.

In our thesis we have addressed this by letting the algorithms explore in *computer simulations* of industrial processes. The simulations were computer programs designed to behave like real chemical plants and a truck brake. By doing so we essentially only change the type of the game to become like a real industrial process.

We asked the reinforcement learning algorithms to play the simulations in a way such that the outgoing signals from the chemical plants and the brake would follow setpoints. The algorithms succeeded to do so but it took tens of thousands of failed "games" to learn.

In the current reinforcement learning algorithms, we have very little knowledge of exactly what it is going to do when playing a game. This, together with the need for many failures to learn, must be overcome before deploying such algorithms into the industry.

One promising current research is about how reinforcement learning algorithms can be connected to classical automatic control techniques. This would allow for the stability of established control techniques *and* let the reinforcement learning algorithm assist with the part of the problem where classic control has trouble.

Some parts of the reinforcement learning algorithms will, with a high chance, nestle their way into the industry in the coming years. This will probably increase efficiency, safety and help automate the industry. Thus we are happy to help shed a ray of light in this void. It has been a very enjoyable journey which we encourage everyone to join in a way they seem fit.

To infinity and beyond!..

References

[1] David Silver Demi Hassabis. Alphago zero: Starting from scratch, 2017.

[2] Peter Norvig. On the (small) number of atoms in the universe, 2016.