Predicting the Future Using Deep Learning

Magnus Wallgren

April 18, 2019

Deep learning methods can teach computers to make short-term predictions of the future. In this work a specific method for training has been tested to see how it performs on prediction problems.

The possibilities that come with the ability to teach computers to predict the future are many. One example is autonomous vehicles. Making autonomous vehicles such as self driving cars involves many challenges. Perhaps the most challenging aspect is that the vehicles must be equipped with the ability to understand complicated traffic situations, anticipate what will happen, and then act accordingly. This means that the cars must be designed to be good at predicting the future to overcome this challenge.

Another interesting application is human-robot interaction. In eye-to-eye conversations between people, speech is not the only way that people communicate. Body language, facial expressions and other visual clues also play a part in the conversation. In particular we predict the different visual queues and mimic or otherwise adapt to what other people are doing. If robots are going to have natural communication with humans, then they must able to read the body language and facial expressions in the same way as people do.

Deep learning is a method for teaching computers relationships in data using an artificial neural network. Artificial neural networks are inspired by the neural networks in the human brain and use a mathematical model of the neurons to learn relathinships between data. When using deep learning particular care has to be taken to choose something called loss function. The loss function is a way for the network to know what errors it is making and learn from them. The simple loss functions compare the output of the network to the actual data and construct an error from that. In prediction problems there can be possible future outcomes that are not represented in the data, which the simple loss functions do not take into account. The network might make one possible prediction, but one that is different from the one represented in the data. Generative Adversarial Networks (GAN) can be used to mitigate this problem.

Generative Adversarial Networks is a method where, instead of using traditional loss funtions, another network is used to represent the error of the first network. The two networks learn in parallel with adverse objectives. The goal for the adverse network, usually called the discriminator, is to distinguish between the predictions made by the generator network and the real data samples, while the generator network is learning to make predictions that can fool the discriminator network as real samples. It is a game between the networks and when the discriminator network improves, so must also the generator network improve to keep up.

We cannot make computers predict when we will be driven by autonomous vehicles and being served by humanoid robots in restaurants, but they can play a role in making it happen.