BUSINESS & INFORMATION TECHNOLOGY (BIT) CLUB NEWSLETTER

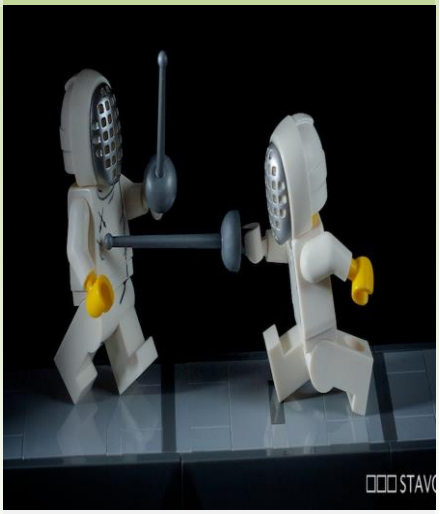

Contents

DUELING NEURAL NETWORKS

Artificial intelligence is getting very good at identifying things : show it a million pictures, and it can tell you with uncanny accuracy which ones depict a pedestrian crossing a street. But AI is hopeless at generating images of pedestrians by itself. If it could do that, it would be able to create gobs of realistic but synthetic pictures depicting pedestrians in various settings, which a self-driving car could use to train itself without ever going out on the road.

The problem is, creating something entirely new requires imagination—and until now that has perplexed AI's.

The solution first occurred to **Ian Goodfellow**, a PhD student at the University of Montreal, during an academic argument in a bar in 2014. The approach, known as a Generative Adversarial Network (GAN), takes two neural networks—the simplified mathematical models of the human brain that underpin most modern machine learning—and pits them against each other in a digital cat-and-mouse game.

Ian Goodfellow

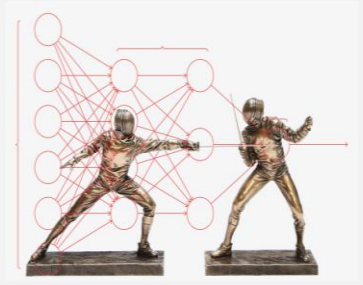

Getting GAN'S to work well can be tricky. If there are glitches, the results can be bizarre.

Both the networks are trained on the same data set. One, known as the 'generator', is tasked with creating variations on images it's already seen—perhaps a picture of a pedestrian with an extra arm. The second, known as the 'discriminator', is asked to identify whether the example it sees is like the images it has been trained on or if it's a fake produced by the generator—basically, is that three-armed person likely to be real or not?

Over time, the 'generator' can become so good at producing images that the 'discriminator' won't be able to spot the fakes. Essentially, the 'generator' has been taught to recognize, and then create, realistic-looking images of pedestrians.

Why It Matters?

Dueling Neural Networks gives machines something akin to a sense of imagination, which may help them become less reliant on humans—but also turns them into alarmingly powerful tools for digital fakery.

The results aren't always perfect: GAN's can conjure up bicycles with two sets of handlebars, say, or faces with eyebrows in the wrong place. But because the images and sounds are often startlingly realistic, some experts believe there's a sense in which GAN's are beginning to understand the underlying structure of the world they see and hear. And that means AI may gain, along with a sense of imagination, a more independent ability to make sense of what it sees in the world.

—Jamie Condliffe

The technology has become one of the most promising advances in AI in the past decade, able to help machines produce results that fool even humans.

The breakthrough is two AI systems sparring with each other to create ultra-realistic original images or sounds, something machines have never been able to do before.

GAN's have been put to use creating realistic-sounding speech and photorealistic fake imagery. In one example, researchers from chipmaker Nvidia primed a GAN with celebrity photographs to create hundreds of credible faces of people who don't exist. Another research group made not-unconvincing fake paintings that look like the works of Van Gogh. Pushed further, GAN's can reimagine images in different ways—making a sunny road appear snowy, or turning horses into zebras.

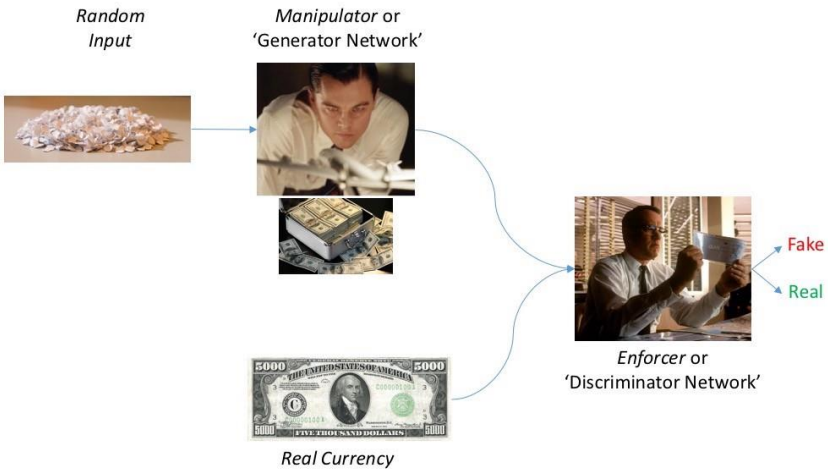

Dueling Neural Networks or Generative Adversarial Network (GAN)

Now let's get our story and actors translated, in the context of GAN's. Both the *manipulator* and *enforcer* are models, variants of the Deep learning neural networks.

The *manipulator* is called the *'Generator network'* which is tasked with the job of creating training data, starting randomly and getting as realistic as it can. The *enforcer* is the *'Discriminator network'*, whose job is to detect and classify these as real or fake, and become pretty good at it.

By pairing two models against each other as adversaries, we set them up for a healthy competition. Each trying to master its own job across thousands of iterations, with no manual intervention. And voila, we end up with true-looking fakes and also a model that can detect most con-jobs.

And this is why GAN's are such a master stroke in AI since they solve both the real-world problems of generating data when you don't have enough to start with, and training models with no manual intervention, a form of unsupervised learning.

Atleast that's where they are headed, and they are already operational. Over the past couple of years, there have been steady advancement of GAN's with hundreds of variants created, and many more innovations underway.

Setting off the perfect cat-and-mouse game

The Dueling Neural Networks was designed based on the following theory. Imagine a movie where two estranged brothers embrace opposing philosophies in life. One starts a fresh underworld operation of printing fake currencies as a 'manipulator', and the other enrolls in a bureau to set up a new division that detects counterfeits as an 'enforcer'.

Let's say that the 'manipulator' in underworld starts with a disadvantage of knowing nothing about what original currencies look like. The 'enforcer' in the bureau knows just basics of how real currencies look. And then the game begins.

The 'manipulator' starts printing, but the initial fakes are terrible. It doesn't need even a trained eye to detect the counterfeits and promptly every single one of them is detected by the 'enforcer'.

The 'manipulator' is industrious and keeps churning out fakes, while also learning what didn't work in previous attempts. By sheer magnitude of experimentation with fakes & some feedback, the quality of counterfeits slowly starts inching up.

Eventually, the 'manipulator' starts getting a few random counterfeits right and this goes undetected by the 'enforcer'. So, its learning time on the other side and the 'enforcer' takes lessons on detecting these smarter counterfeits. With the 'enforcer' getting smarter, the counterfeits are detected again. The 'manipulator' has no choice but to upgrade the counterfeiting operation to create more genuine-looking fakes.

This continuous game of cat-and-mouse continues and ends up making experts out of both the 'manipulator' and 'enforcer' so much that the counterfeits are indistinguishable from the genuine ones, and also the detection of such ingenious fakes becomes almost uncanny.

Utility of Dueling Neural Networks

Generation of training data

GAN's do the heavy lifting of creating tons of training data, which can shift AI into the fast-lane of progress. Imagine GAN's spawning realistic 3D worlds similar to ours, with millions of miles of roads & all possible traffic scenarios.
Rather than a self-driving car or drone getting trained in the real-world and causing horrendous accidents, they could get trained in these virtual worlds and become expert drivers. With GPU computing, this can be instantaneous.
While these are directional applications, GAN's have already been applied to high-impact business applications like drug discovery and there are literally hundreds of use cases in early stages of experimentation.
While this might already sound revolutionary, the best in GAN's is yet to come.

Key Players of GAN

Google Brain, DeepMind, Nvidia

1. Creative pursuits

It's incredible to imagine that machines have finally unlocked their right brains. After all, who wouldn't be surprised when a nerdy programmer suddenly starts penning award-winning poetry.

With a new-found approach to mimic real images, GAN's have started creating imaginary celebrities or new masterpieces that bear distinctive signature of artists. The potential use cases with this ability spans creative disciplines.

2. Translating text

Suppose you want to find out how a person would look without their glasses or with a new hairdo, you just have to ask to have it created. Not very different from asking for the day's weather or mapping your upcoming commute.

By creating new flora and fauna to user specification of a short description, GAN's have been dutifully granting the demands, just like a wish-fulfilling genie. Pity they couldn't breathe life into the creations, atleast not yet.