Neural Networks

Two AI systems can compete with each other to **create** ultra-realistic original images or sounds, something machines have never been able to do before. This gives machines something similar 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.

Artificial intelligence is getting very good at identifying things: show it a million pictures, and it can **accurately** tell you 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 a large amount 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 that creating something **entirely** new **requires** imagination—and until now that has confused AIs. The **solution** first **occurred** to Ian Goodfellow, then 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, or GAN, takes two neural **networks**—the **simplified** mathematical models of the human brain that support most modern machine learning—and sets them against each other in a digital cat-and-mouse game.

Both networks are trained on the same data set. One, known as the generator, is tasked with creating **variations** on images it has 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 a fake produced by the generator—**basically**, is that three-armed person likely to be real?

Over time, the generator can become so good at producing images that the discriminator cannot identify fakes. **Essentially**, the generator has been taught to **recognize** and create realistic-looking images of pedestrians. GANs have been put to use creating realistic-sounding **speech** and photorealistic fake **imagery**.

In one example, **researchers** from chipmaker Nvidia configured a GAN with celebrity photographs to create hundreds of credible faces of people who don't exist. Another research group made fake paintings that look like the works of van Gogh. Pushed further, GANs can reimagine images in different ways—making a sunny road **appear** snowy, or turning horses into zebras.

The **results** aren't always perfect: GANs can generate bicycles with two **sets** of handlebars, or faces with eyebrows in the wrong place. But because the images and sounds are often surprisingly realistic, some experts believe there is a sense in which GANs 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.

Materials' Quantum Leap

IBM has simulated the electronic structure of a small molecule, using a seven-qubit quantum computer

Understanding molecules in **exact** detail will **allow** chemists to design more effective drugs and better materials for generating and **distributing** energy.

The prospect of powerful new quantum computers comes with a **puzzle**. They will be capable of feats of computation **inconceivable** with today's machines, but we haven't yet figured out what we might do with those powers. One **likely** and enticing possibility: precisely designing molecules.

Chemists are already dreaming of new proteins for far more effective drugs, **novel** electrolytes for better batteries, **compounds** that could turn sunlight directly into a liquid fuel, and much more efficient solar **cells**. We don't have these things because molecules are hard to model on a classical computer. Try simulating the **behavior** of the electrons in even a relatively simple molecule and you **run into** complexities far beyond the **capabilities** of today's computers.

But it is a natural problem for quantum computers, which instead of digital bits representing 1s and 0s use "qubits" that are themselves quantum systems. **Recently**, IBM researchers used a quantum computer with seven qubits to model a small molecule made of three atoms. It should become possible to accurately simulate far larger and more interesting molecules as scientists build machines with more qubits and, just as important, better quantum algorithms.

Adapted from **MIT Technology Review**