Generative Adversarial Networks (GANs)

"We propose a new framework for estimating generative models via an adversarial process, in which we simultaneously train two models: a generative model G that captures the data distribution, and a discriminative model D that estimates the probability that a sample came from the training data rather than G. The training procedure for G is to maximize the probability of D making a mistake. This framework corresponds to a minimax two-player game. In the space of arbitrary functions G and D, a unique solution exists, with G recovering the training data distribution and D equal to ½ everywhere. In the case where G and D are defined by multilayer perceptrons, the entire system can be trained with backpropagation." - Goodfellow et. al.

Generative Adversarial Nets, Retrieved from:
https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf

---

If you follow technology news, you might have seen a website called "ThisPersonDoesNotExist.com". When you visit this website, it gives you an image of a random person's face. However, these faces are actually generated by a computer program - the person in the image does not really exist. This computer program works by using something called a *Generative Adversarial Network* (GAN).

A GAN is actually made up of two sub-programs, a *generator* and a *discriminator*. The generator's job is to try and generate images that look like real human face. The *discriminators* job is to *discriminate* (or decide) whether any given face is real or fake. A common analogy for GANs is to think of it as an art countifiter (the generator) versus an art inspector (the discriminator).

At first, the counterfeiter does not know what the difference is between fake artwork and a real artwork is, so it produces images at random. As time goes on, the counterfeiter will learn what makes a good fake or not. It learns based off of whether the art inspector thought the counterfeit was real or not. Meanwhile, the art inspector becomes aware of the counterfeiters tricks, and becomes more careful in their observations.

This game forms an ever-improving feedback loop, which leads the counterfeiter to generate more and more realistic fakes. Eventually, the discriminator can no longer tell real from fake. This is the point where we can generate extremely realistic faces like the website above.