

Generative adversarial network

A generative adversarial network (GAN) is a class of machine learning frameworks invented by Ian Goodfellow and his colleagues in 2014.[1] Two neural networks contest with each other in a game (in the sense of game theory, often but not always in the form of a zero-sum game). Given a training set, this technique learns to generate new data with the same statistics as the training set. For example, a GAN trained on photographs can generate new photographs that look at least superficially authentic to human observers, having many realistic characteristics. Though originally proposed as a form of generative model for unsupervised learning, GANs have also proven useful for semi-supervised learning,[2] fully supervised learning,[3] and reinforcement learning.[4] In a 2016 seminar, Yann LeCun described GANs as "the coolest idea in machine learning in the last twenty years".[5]

Method

The generative network generates candidates while the discriminative network evaluates them.[1] The contest operates in terms of data distributions. Typically, the generative network learns to map from a latent space to a data distribution of interest, while the discriminative network distinguishes candidates produced by the generator from the true data distribution. The generative network's training objective is to increase the error rate of the discriminative network (i.e., "fool" the discriminator network by producing novel candidates that the discriminator thinks are not synthesized (are part of the true data distribution)).[1][6]

A known dataset serves as the initial training data for the discriminator. Training it involves presenting it with samples from the training dataset, until it achieves acceptable accuracy. The generator trains based on whether it succeeds in fooling the discriminator. Typically, the generator is seeded with randomized input that is sampled from a predefined latent space (e.g. a multivariate normal distribution). Thereafter, candidates synthesized by the generator are evaluated by the discriminator. Backpropagation is applied in both networks so that the generator produces better images, while the discriminator becomes more skilled at flagging synthetic images.[7] The generator is typically a deconvolutional neural network, and the discriminator is a convolutional neural network.

GANs often suffer from a "mode collapse" where they fail to generalize properly, missing entire modes from the input data. For example, a GAN trained on the MNIST dataset containing many samples of each digit, might nevertheless timidly omit a subset of the digits from its output. Some researchers perceive the root problem to be a weak discriminative network that fails to notice the pattern of omission, while others assign blame to a bad choice of objective function. Many solutions have been proposed.[8]

Applications

GAN applications have increased rapidly.[9]

Fashion, art and advertising

GANs can be used to create photos of imaginary fashion models, with no need to hire a model, photographer, makeup artist, or pay for a studio and transportation.[10] GANs can be used to create fashion advertising campaigns including more diverse groups of models, which may increase

intent to buy among people resembling the models. [11] GANs can also be used to create portraits, landscapes, and album covers (<https://www.artbreeder.com/>).

Science

GANs can improve astronomical images [12] and simulate gravitational lensing for dark matter research. [13][14][15] They were used in 2019 to successfully model the distribution of dark matter in a particular direction in space and to predict the gravitational lensing that will occur. [16][17]

GANs have been proposed as a fast and accurate way of modelling high energy jet formation [18] and modelling showers through calorimeters of high-energy physics experiments. [19][20][21][22] GANs have also been trained to accurately approximate bottlenecks in computationally expensive simulations of particle physics experiments. Applications in the context of present and proposed CERN experiments have demonstrated the potential of these methods for accelerating simulation and/or improving simulation fidelity. [23][24]

Video games

In 2018, GANs reached the video game modding community, as a method of up-scaling low-resolution 2D textures in old video games by recreating them in 4k or higher resolutions via image training, and then down-sampling them to fit the game's native resolution (with results resembling the supersampling method of anti-aliasing). [25] With proper training, GANs provide a clearer and sharper 2D texture image magnitudes higher in quality than the original, while fully retaining the original's level of details, colors, etc. Known examples of extensive GAN usage include Final Fantasy VIII, Final Fantasy IX, Resident Evil Remake HD Remaster, and Max Payne.

Concerns about malicious applications

Concerns have been raised about the potential use of GAN-based human image synthesis for sinister purposes, e.g., to produce fake, possibly incriminating, photographs and videos. [26] GANs can be used to generate unique, realistic profile photos of people who do not exist, in order to automate creation of fake social media profiles. [27]

In 2019 the state of California considered [28] and passed on October 3, 2019 the bill AB-602 (https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201920200AB602), which bans the use of human image synthesis technologies to make fake pornography without the consent of the people depicted, and bill AB-730

(https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201920200AB730), which prohibits distribution of manipulated videos of a political candidate within 60 days of an election. Both bills were authored by Assembly member Marc Berman and signed by Governor Gavin Newsom. The laws will come into effect in 2020. [29]

DARPA's Media Forensics program studies ways to counteract fake media, including fake media produced using GANs. [30]

Miscellaneous applications

GAN can be used to detect glaucomatous images helping the early diagnosis which is essential to avoid partial or total loss of vision [31].

GANs that produce photorealistic images can be used to visualize interior design, industrial design, shoes,[32] bags, and clothing items or items for computer games' scenes. Such networks were reported to be used by Facebook. [33]

GANs can reconstruct 3D models of objects from images, [34] and model patterns of motion in video. [35]

GANs can be used to age face photographs to show how an individual's appearance might change with age. [36]

GANs can also be used to transfer map styles in cartography [37] or augment street view imagery. [38]

A variation of the GANs is used in training a network to generate optimal control inputs to nonlinear dynamical systems. Where the discriminatory network is known as a critic that checks the optimality of the solution and the generative network is known as an Adaptive network that generates the optimal control. The critic and adaptive network train each other to approximate a nonlinear optimal control. [39]

GANs have been used to visualize the effect that climate change will have on specific houses. [40]

A GAN model called Speech2Face can reconstruct an image of a person's face after listening to their voice. [41]

In 2016 GANs were used to generate new molecules for a variety of protein targets implicated in cancer, inflammation, and fibrosis. In 2019 GAN-generated molecules were validated experimentally all the way into mice. [42][43]

History

The most direct inspiration for GANs was noise-contrastive estimation, [44] which uses the same loss function as GANs and which Goodfellow studied during his PhD in 2010–2014.

Other people had similar ideas but did not develop them similarly. An idea involving adversarial networks was published in a 2010 blog post by Olli Niemitalo. [45] This idea was never implemented and did not involve stochasticity in the generator and thus was not a generative model. It is now known as a conditional GAN or cGAN. [46] An idea similar to GANs was used to model animal behavior by Li, Gauci and Gross in 2013. [47]

Adversarial machine learning has other uses besides generative modeling and can be applied to models other than neural networks. In control theory, adversarial learning based on neural networks was used in 2006 to train robust controllers in a game theoretic sense, by alternating the iterations between a minimizer policy, the controller, and a maximizer policy, the disturbance [48][49].

In 2017, a GAN was used for image enhancement focusing on realistic textures rather than pixel-accuracy, producing a higher image quality at high magnification. [50] In 2017, the first faces were generated. [51] These were exhibited in February 2018 at the Grand Palais. [52][53] Faces generated by StyleGAN [54] in 2019 drew comparisons with deepfakes. [55][56][57]

Beginning in 2017, GAN technology began to make its presence felt in the fine arts arena with the appearance of a newly developed implementation which was said to have crossed the threshold of being able to generate unique and appealing abstract paintings, and thus dubbed a "CAN", for "creative adversarial network". [58] A GAN system was used to create the 2018 painting Edmond de Belamy, which sold for US \$432,500. [59] An early 2019 article by members of the original CAN team discussed further progress with that system, and gave consideration as well to the overall prospects for an AI-enabled art. [60]

In May 2019, researchers at Samsung demonstrated a GAN-based system that produces videos of a person speaking given only a single photo of that person. [61]

In August 2019, A large dataset consisting of 12,197 MIDI songs each with paired lyrics and melody alignment was created for neural melody generation from lyrics using conditional GAN-LSTM (refer to sources at GitHub AI Melody Generation from Lyrics (<https://github.com/yy1lab/Lyrics-Conditioned-Neural-Melody-Generation>)). [62]

Classification

Bidirectional GAN

Bidirectional GAN (BiGAN) aims to introduce a generator model to act as the discriminator, whereby the discriminator naturally considers the entire translation space so that the inadequate training problem can be alleviated. To satisfy this property, generator and discriminator are both designed to model the joint probability of sentence pairs, with the difference that, the generator decomposes the joint probability with a source language model and a source-to-target translation model, while the discriminator is formulated as a target language model and a target-to-source translation model. To further leverage the symmetry of them, an auxiliary GAN is introduced and adopts generator and discriminator models of original one as its own discriminator and generator respectively. Two GANs are alternately trained to update the parameters. The resulting learned feature representation is useful for auxiliary supervised discrimination tasks, competitive with contemporary approaches to unsupervised and self-supervised feature learning. [63]

--For the references look up Wikipedia --