Chapter 2

Generative Machine Learning

Charlie Harper, PhD
Case Western Reserve University

Introduction

Generative machine learning is a hot topic. With the 2020 election approaching, Facebook and Reddit have each issued their own bans on the category of machine-generated or -altered content that is commonly termed "deep fakes" (Cohen 2020; Romm, Harwell, and Stanley-Becker 2020). Calls for regulation of the broader, and very nebulous category of fake news are now part of US political debates, too. Although well known and often discussed in newspapers and on TV because of their dystopian implications, deep fakes are just one application of generative machine learning. There is a remarkable need for others, especially humanists and social scientists, to become involved in discussions about the future uses of this technology, but this first requires a broader awareness of generative machine learning's functioning and power. Many articles on the subject of generative machine learning exist in specialized, highly technical literature, but there is little that covers this topic for a broader audience while retaining important high-level information on how the technology actually operates.

This chapter presents an overview of generative machine learning with particular focus on generative adversarial networks (GANs). GANs are largely responsible for the revolution in machine-generated content that has occured in the past few years and their impact on our future extends well beyond that of producing purposefully-deceptive fakes. After covering generative learning and the working of GANs, this chapter touches on some interesting and significant applications of GANs that are not likely to be familiar to the reader. The hope is that this will serve as the start of a larger discussion on generative learning outside of the confines of technical literature and sensational news stories.
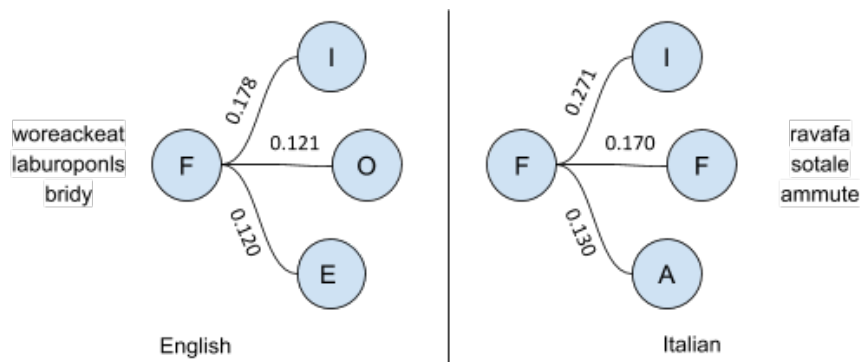
Figure 2.1: The three most-common letters following "F" in two Markov chains trained on an English and Italian dictionary. Three examples of generated words are given for each Markov chain that show how the Markov chain captures high-level information about letter arrangements in the different languages.

What is Generative Machine Learning?

Machine learning, which is a subdomain of Artificial Intelligence, is roughly divided into three paradigms that rely on different methods of learning: supervised, unsupervised, and reinforcement learning (Murphy 2012, 1–15; Burkov 2019, 1–8). These differ in the types of datasets used for learning and the desired applications. Supervised and unsupervised machine learning use labeled and unlabeled datasets, respectively, to assign unseen data to human-generated labels or statistically-constructed groups. Both supervised and unsupervised approaches are commonly used for classification and regression problems, where we wish to predict categorical or quantitative information about new data. A combined form of these two paradigms, called semi-supervised learning, that mixes labeled and unlabeled data also exists. Reinforcement learning, on the other hand, is a paradigm in which an agent learns how to function in a specific environment by being rewarded or penalized for its behavior. For example, reinforcement learning can be used to train a robot to successfully navigate around obstacles in a physical space.

Generative machine learning, rather than being a specific learning paradigm, encompasses an ever-growing variety of techniques that are capable of generating new data based on learned patterns. The process of learning these patterns can engage both supervised and unsupervised learning. A simple, statistical example of one type of generative learning is a Markov chain. From a given set of data, a Markov chain calculates and stores the probabilities of a following state based on a current state. For example, a Markov chain can be trained on a list of English words to store the probabilities of any one letter occuring after another letter. These probabilities chain together to represent that chance of moving from the current letter state (e.g. the letter q) to a succeeding letter state (e.g. the letter u) based on the data from which it has learned.

If another Markov chain were trained on Italian words instead of English, the probabilities would change, and for this reason, Markov chains can capture important high level information about datasets (Figure 2.1). They can then be sampled to generate new data by starting from a random state and probabilistically moving to succeeding states. In figure 2.1, you can see the
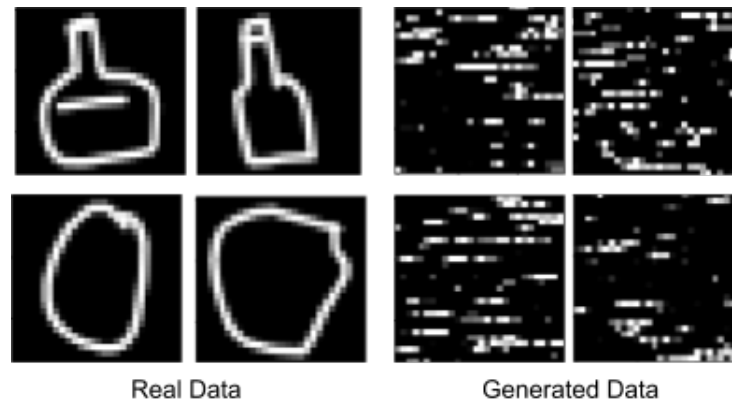
Real Data   Generated Data

Figure 2.2: Images generated with a simple statistical model appear as noise as the model is insufficient to capture the structure of the real data (Markov chains trained using wine bottles and circles from Google's QuickDraw dataset).

probability that the letter "F" transitions to the three most common succeeding letters in English and Italian. A few examples of "words" generated by two Markov chains trained on an English and Italian dictionary are also given. The example words are generated by sampling the probability distributions of the Markov chain, letter by letter, so that the generated words are statistically random, but guided by the learned probability of one letter following another. The different probabilities of letter combinations in English and Italian result in distinctly different generated words. This exemplifies how a generative model can capture specific aspects of a dataset to create new data.

The letter combinations are nonsense, but they still reflect the high-level structure of Italian and English words in the way letters join together, such as the different utilization of vowels in each language. These basic Markov chains demonstrate the essence of generative learning: a generative approach learns a distribution over a dataset, or in other words, a mathematical representation of a dataset, which can then be sampled to generate new data that exists within the learned structure of that dataset. How convincing the generated data appears to a human observer depends on the type and tuning of the machine learning model chosen and the data upon which the model has been trained. So, what happens if we build a comparable Markov chain with image data[1] instead of words, and then sample, pixel by pixel, from it to generate new images? The results are just noise and the generated images reveal no hint of a wine bottle or circle to the human eye (Figure 2.2).

The very simple generative statistical model we have chosen to use is incapable of capturing the distribution of the underlying images sufficiently enough to produce realistic new images. Other types of generative statistical models, like Naive Bayes or a higher-order Markov chain,[2]

---

[1] In many examples, I have used the Google QuickDraw Dataset to highlight features of generative machine learning. The dataset is freely available (`https://github.com/googlecreativelab/quickdraw-dataset`) and licensed under CC BY 4.0.

[2] The order of a Markov chain reflects how many preceding states are taken into account. For example, a 2nd order Markov chain would look at the preceding two letters to calculate the probability of a succeeding letter. Rudimentary autocomplete is a good example of Markov chains in application.

could perhaps capture a bit more information about the training data, but they would still be insufficient for real-world applications like this.[3] Image, video, and audio are complicated; it is hard to reduce them to their essence with basic statistical rules in the way we were able to with the ordering of letters in English and Italian. Capturing the intricate and often-inscrutable distributions that underlie real-world media, like full-sized photographs of people, is where deep (i.e. using neural networks) generative learning shines and where generative adversarial networks have revolutionized machine-generated content.

Generative Adversarial Networks

The problem of capturing the complexity of an image so that a computer can generate new images leads directly to the emergence of Generative Adversarial Networks, which are a neural-network-based model architecture within the broader sphere of generative machine learning. Although prior deep learning approaches to generating data, particularly variational autoencoders, already existed, it was a breakthrough in 2014 that changed the fabric and power of generative machine learning. Like every big development, it has an origin story that has moved into legend with its many retellings. According to the handed-down tale (Giles 2018), in 2014 doctoral student Ian Goodfellow was at a bar with friends when the topic of generating photos arose. His friends were working out a method to create realistic images by using complex statistical analyses of existing images. Goodfellow countered that it would not work; there were too many variables at play within such data. Instead, he put forth the idea of pairing two neural networks against each other in a type of zero-sum game where the goal was to generate believable fake images. According to the story, he developed this idea into working code that night and his paired neural network architecture produced results the very first time. This was the birth of Generative Adversarial Networks or GANs. Goodfellow's work was quickly disseminated in what is one of the most influential papers in the recent history of machine learning (Goodfellow et al. 2014).

GANs have progressed in almost miraculous ways since 2014, but the crux of their architecture remains the coupling of two neural networks. Each neural network has a specific function in the pairing. The first network, called the generator, is tasked with generating fake examples of some dataset. To produce this data it randomly samples from an n-dimensional latent space often labeled Z. In simple terms, the generator takes random noise (really a random list of n-numbers where n is the dimensionality of the latent space) as its input and outputs its attempt at a fake piece of data, such as an image, clip of audio, or row of tabular information. The second neural network, called the discriminator, takes both fake and real data as input. Its role is to correctly discriminate between fake and real examples.[4] The generator and discriminator networks are then coupled together as adversaries, hence "adversarial" in the name. The output from the generator flows into the discriminator, and information on the success or failure of the discriminator to identify fakes (i.e. the discriminator's loss) flows back through the network so that the generator and discriminator each knows how well it is performing compared to the other. All of this happens automatically, without any need for human supervision. When the generator finds it is doing poorly, it learns to produce better examples by updating its weights and biases through traditional backpropagation (see especially Langr and Bok 2019, 3–16 for a more detailed summary of this). As backpropagation updates the generator network's weights and biases, the generator

---

[3]This is not to imply that these models do not have immense practical applications in other areas of machine learning.
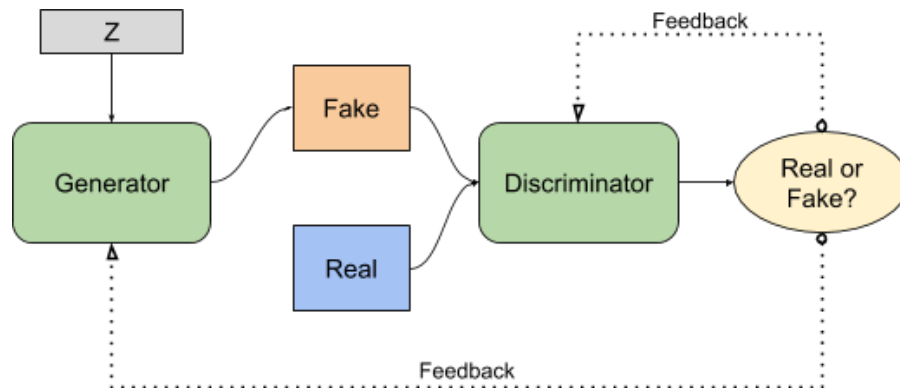[4]Its function is exactly that of any other binary classifier found in machine learning.

Figure 2.3: At the heart of a GAN are two neural networks, the generator and the discriminator. As the generator learns to produce fake data, the discriminator learns to separate it out. The pairing of the two in an adversarial structure forces each to improve at its given task.

Figure 2.4: A GAN being trained on wine bottle sketches from Google's quickdraw dataset (`https://github.com/googlecreativelab/quickdraw-dataset`) shows the generator learning how to produce better sketches over time. Moving from left to right, the generator begins by outputting random noise and progressively generates better sketches as it tries to trick the discriminator.

inherently begins to map regions of the randomly sampled Z space to characteristics found in the real dataset. Contrarily, as the discriminator finds that it is not identifying better fakes accurately, it learns to separate these out in new ways.

At first, the generator outputs random data and the discriminator easily catches these fakes (Figure 2.4). As the results of the discriminator feed back into the generator, however, the generator learns to trick its foe by creating more convincing fakes. The discriminator consecutively learns to better separate out these more convincing fakes. Turn after turn, the two networks drive one another to become better at their specialized tasks and the generated data becomes increasingly like the real data.[5] At the end of training, ideally, it will not be possible to distinguish between real and fake (Figure 2.5).

In the original publication, the first GANs were trained on sets of small images, like the Toronto Face Dataset, which contains 32×32 pixel grayscale photos of faces and facial expressions (Goodfellow et al. 2014). Although the generator's results were convincing when compared to the originals, the fake images were still small, colorless, and pixelated. Since then an explosion of research into GANs and increased computational power has led to strikingly realis-

---

[5]See `https://poloclub.github.io/ganlab/` (accessed Jan 17, 2020) (Kahng et al. 2019).
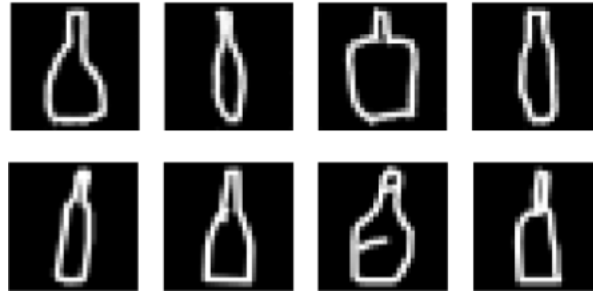
Figure 2.5: The fully trained generator from Figure 2.4 produces examples that are not readily distinguishable from real world data. The top row of sketches were produced by the GAN and the bottom row were drawn by humans.

tic images. The most recent milestone was reached in 2019 by researchers with NVIDIA, who built a GAN that generates high-quality photo-realistic images of people (Karras, Laine, and Aila 2019). When contrasted with the results of 2014 (Figure 2.6), the stunning progression of GANs is self-evident, and it is difficult to believe that the person on the right does not exist.

Some Applications of Generative Adversarial Networks

Over the past five years, many papers on implementations of GANs have been released by researchers (Alqahtani, Kavakli-Thorne, and Kumar 2019; Wang, She, and Ward 2019). The list of applications is extensive and ever growing, but it is worth pointing out some of the major examples as of 2019 and why they are significant. These examples highlight the vast power of GANs and underscore the importance of understanding and carefully scrutinizing this type of machine learning.

Data Augmentation

One major problem in machine learning has always been the lack of labeled datasets, which are required by supervised learning approaches. Labeling data is time consuming and expensive. Without good labeled data, trained models are limited in their power to learn and in their ability to generalize to real-world problems. Services, such as Amazon's Mechanical Turk, have attempted to crowdsource the tedious process of manually assigning labels to data, but labeling has remained a bottleneck in machine learning. GANs are helping to alleviate this bottleneck by generating new labeled data that is indistinguishable from the real data. This process can grow a small labeled dataset into one that is larger and more useful for training purposes. In the area of medical imaging and diagnostics this may have profound effects (Yi, Walia, and Babyn 2019). For example, GANs can produce photorealistic images of skin lesions that expert dermatologists are able to separate from real images only slightly over 50% of the time (Baur, Albarqouni, and Navab 2018) and they can synthesize high-resolution mammograms for training better cancer detection algorithms (Korkinof et al. 2018).

A corollary effect of these developments in medical imaging is the potential to publicly release

2014
(Goodfellow et al. 2014, fig. 2b)

2019
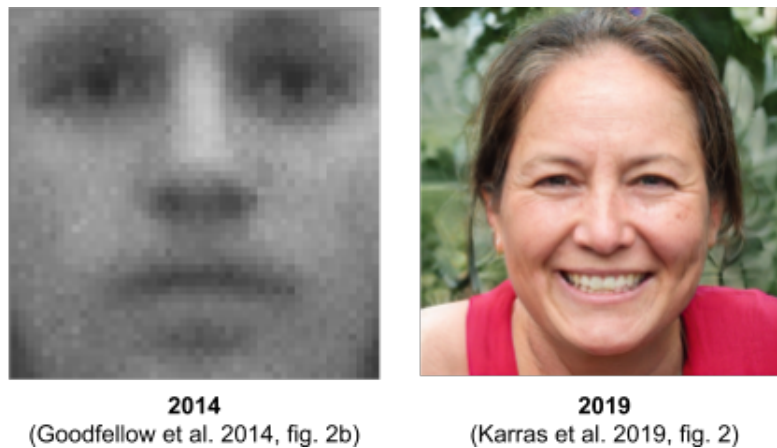(Karras et al. 2019, fig. 2)

Figure 2.6: An image of a generated face from the original GAN publication (left) and the 2019 milestone (right) shows how the ability of GANs to produce photo-realistic images has evolved since 2014.

large medical datasets and thereby expand researchers' access to important data. Whereas the dissemination of traditional medical images is constrained by strict health privacy laws, generated images may not be governed by such rules. I qualify this statement with "may", because any restrictions or ethical guidelines for the use of medical data that is generated from real patient data requires extensive discussion and legal reviews that have not yet happened. Under certain conditions, it may also be possible to infer original data from a GAN (Mukherjee et al. 2019). How institutional review boards, professional medical organizations, and courts weigh in on this topic will be seen in the coming years.

In addition to generating entirely new data, a GAN can augment datasets by expanding their coverage to new domains. For example, autonomous vehicles must cope with an array of road and weather conditions that are unpredictable. Training a model to identify pedestrians, street signs, road lines, and so on with images taken on a sunny day will not translate well to variable real-world conditions. Using one dataset, in a process known as style transfer, GANs can translate one image to other domains (Figure 2.7). This can include creating night road scenes from day scenes (Romera et al. 2019) and producing images of street signs under varying lighting conditions (Chowdhury et al. 2019). This added data permits models to account for greater variability under operating conditions without the high cost of photographing all possible conditions and manually labeling them. Beyond medicine and autonomous vehicles, generative data augmentation will progressively impact other imaging-heavy fields (Shorten and Khoshgoftaar 2019) like remote sensing (L. Ma et al. 2019; D. Ma, Tang, and Zhao 2019).

Creativity and Design

The question of whether machines can possess creativity or artistic ability is philosophically difficult to answer (Mazzone and Elgammal 2019; McCormack, Gifford, and Hutchings 2019). Still, in 2018, Christie's auctioned off its first piece of GAN art for $432,500 (Cohn 2018) and GANs
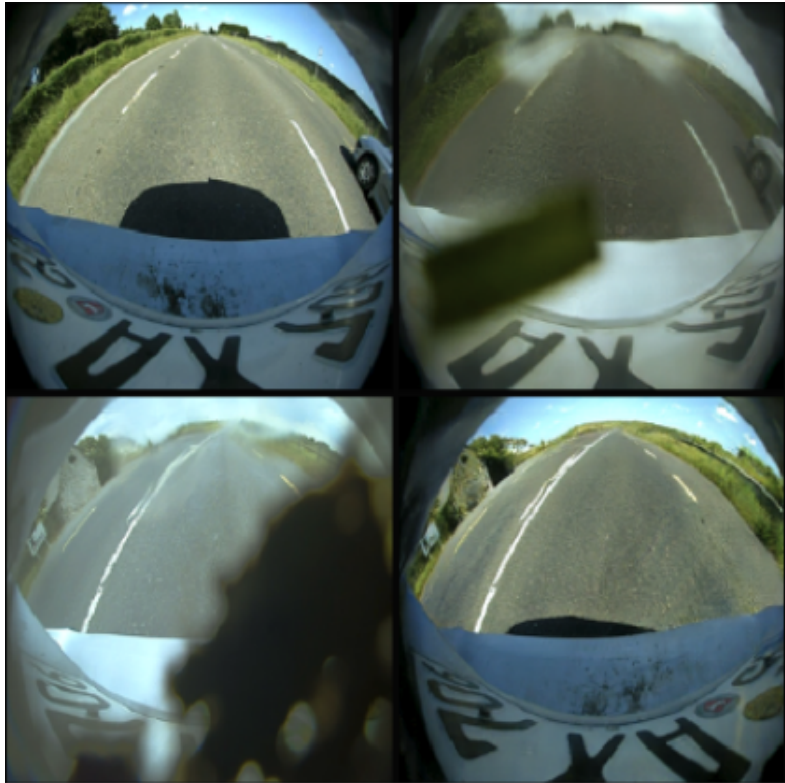
Figure 2.7: The images on the left are originals and the images on the right have been modified by a GAN with the ability to translate images between the domains of "dirty lens" and "clean lens" on a vehicle (from Uřičář et al. 2019, fig. 11).
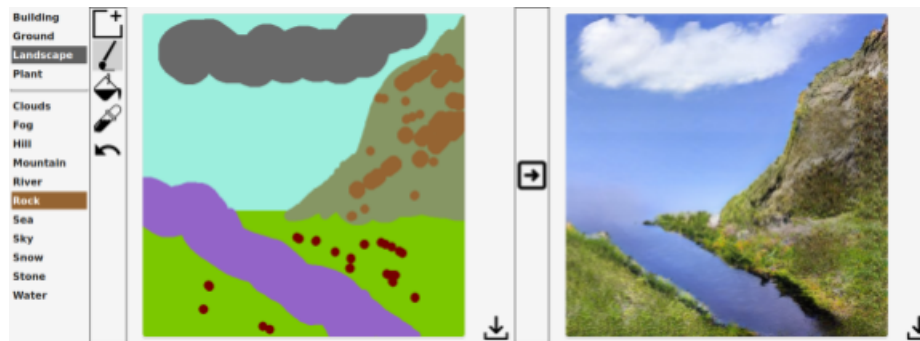
Figure 2.8: This example of GauGAN in action shows a sketched out scene on the left turned into a photo-realistic landscape on the right. *If any representatives of Christie's are reading, the author would be happy to auction this piece.

are increasingly assisting humans in the creative process for all forms of media. Simple models, like CycleGAN, are already able to stylize images in the manner of Van Gogh or Monet (Zhu et al. 2017), and more varied stylistic GANs are emerging.

GauGAN, a beta tool released by NVIDIA, is a great example of GAN-assisted creativity in action. GauGAN allows you to rough out a scene using a paint brush for different categories, like clouds, flowers, and houses (Figure 2.8). It then converts this into a photo reflecting what you have drawn. The online demo[6] remains limited, but the underlying model is powerful and has massive potential (Park et al. 2019). Recently, Martin Scorsese's *The Irishman* made headlines for its digital de-aging of Robert Deniro and other actors. Although this process did not involve GANs, it is highly likely that in the future, GANs will become a major part of cinematic post-production (Giardina 2019) through assistive tools like GauGAN.

Fashion and product design are also being impacted by the use of GANs. Text-to-image synthesis, which can take free text or categories as input to generate a photo-realistic image, has promising potential (Rostamzadeh et al. 2018). By accepting text as input, GANs can let designers rapidly generate new ideas or visualize concepts for products at the start of the design process. For example, a recently published GAN for clothing design accepts basic text and outputs modeled images of the described clothing (Banerjee et al. 2019; Figure 9). In an example of automotive design, a single sketch can be used to generate realistic photos of multiple perspectives of a vehicle (Radhakrishnan et al. 2018). The many fields that rely on quick sketching or visual prototyping, such as architecture or web design, are likely to be influenced by the use of GAN-assisted design software in coming years.

In a similar vein, GANs have an upcoming role in the creation of new medicines, chemicals, and materials (Zhavoronkov 2018). By training a GAN on existing chemical and material structures, research is showing that novel chemicals and materials can be designed with particular properties (Gómez-Bombarelli et al. 2018; Sanchez-Lengeling and Aspuru-Guzik 2018). This is facilitated by how information is encoded in the GAN's latent space (the n-dimensional space from which the generator samples; see "Z" in Figure 2.3). As the generator learns to produce realistic examples, certain aspects of the original data become encoded in regions of the latent

---

[6]See http://nvidia-research-mingyuliu.com/gaugan/ (last accessed January 12, 2019).

Figure 2.9: Text-to-image synthesis can generate images of new fashions based on a description. From the input "maroon round neck mini print a-line bodycon short sleeves" a GAN has produced these three photos (from Banerjee et al. 2019, fig. 11).

Figure 2.10: Two examples of linearly-spaced mappings across the latent space between generated images A and B. Note that by taking one image and moving closer to another, you can alter properties in the image, such as adding steam, removing a cup handle, or changing the angle of view. These characteristics of the dataset are learned by the generator during training and encoded in the latent space. (GAN built on coffee cup sketches from Google's QuickDraw dataset)

space. By moving through this latent space or sampling particular areas, new data with desired properties can then be generated. This can be seen by periodically sampling the latent space and generating an image as one moves between two generated images (Figure 2.10). In the same way, by moving in certain directions or sampling from particular areas of the latent space, new chemicals or medicines with specific properties can be generated.[7]

Impersonation and the Invisible

I have reserved some of the more dystopian and likely more well-heard-of applications of GANs for last. This is the area where GANs' ability to generate convincing media is challenging our perceptions of reality and raising extreme ethical questions (Harper 2018). Deep fakes are, of course, the most well known of these. This can include the creation of fake images, videos, and audio of an individual or the modification of any media to alter what someone appears to be doing or saying. In images and video in particular, GANs make it possible to swap the identity of an individual and manipulate facial attributes or expressions (Tolosana et al. 2020). A large portion

---

[7]This is also relevant to facial manipulation discussed below.

Figure 2.11: GANs are providing a method to reconstruct hidden images of people and objects. Images 1–3 show reconstructions as compared to an input occluded image (OCC) and a ground truth image (GT) (from Fulgeri et al. 2019, fig. 6).

of technical literature is, in fact, now devoted to detecting faked and altered media (see Tolosana et al. 2020, Table IV and V). It remains to be seen how successful any approaches will be. From a theoretical perspective, anything that can detect fakes can also be used to train a better generator since the training process of a GAN is founded on outsmarting a detector (i.e. the discriminator network).

One shocking extension of deep fakes that has emerged is transcript to video creation, which generates a video of someone speaking from a written text. If you want to see this at work, you can view clips of Nixon giving the speech written in the case of an Apollo 11 disaster.[8] As of now, deep fakes like this remain choppy and are largely limited to politicians and celebrities because they require large datasets and additional manipulation, but this limitation is not likely to last. If the evolution of GANs for images is any predictor, the entire emerging field of video generation is likely to progress rapidly. One can imagine the incorporation of text-to-image and deep fakes enabling someone to produce an image of, say, "politican X doing action Y," simply by typing it.

An application of GANs that parallels deep fakes and is likely more menacing in the short term is the infilling or adding of hidden, invisible, or predicted information to existing media. One nascent use is video prediction from an image. For example, in 2017, researchers were able to build a GAN that produced 1-second video clips from a single starting frame (Vondrick and Torralba 2017). This may not seem impressive, but video is notoriously difficult to work with because the content of a succeeding frame can vary so drastically from the preceding frame (for other examples of on-going research into video prediction, see Cai et al. 2018; Wen et al. 2019).

For still images, occluded object reconstruction, in which a GAN is trained to produce a full image of a person or object that is partially hidden behind something else, is progressing (Fulgeri et al. 2019; see Figure 11). For some applications, like autonomous driving, this could save lives as it would help to pick out when a partially-occluded pedestrian is about to emerge from

---

[8]See http://news.mit.edu/2019/mit-apollo-deepfake-art-installation-aims-to-empower-more-discerning-public-1125.

behind a parked car. On the other hand, for surveillance technology, it can further undermine anonymity. Indeed, such GANs are already being explicitly studied for surveillance purposes (Fabbri, Calderara, and Cucchiara 2017). Lastly, I would be remiss if I did not mention that researchers have designed a GAN that can generate an image of what you are thinking about, using EEG signals (Tirupattur et al. 2018).

GANs and the Future

The tension between the creation of more realistic generated data and the technology to detect maliciously generated information is only beginning. The machine learning and data science platform, Kaggle, is replete with publicly-accessible python code for building GANs and detecting fake data. Money, too, is freely flowing in this domain of research; The 2019 Deepfake Detection Challenge sponsored by Facebook, AWS, and Microsoft boasted one million dollars in prizes (`https://www.kaggle.com/c/deepfake-detection-challenge` accessed April 20, 2020). Meanwhile, industry leaders, such as NVidia, continue to fund the training of better and more convincing GANs. The structure of a GAN, with its generator and detector paired adversarially, is now being mirrored in society as groups of researchers competitively work to create and discern generated data. The path that this machine-learning arms race will take is unpredictable, and, therefore, it is all the more important to scrutinize it and make it comprehensible to the broader publics whom it will affect.

References

Alqahtani, Hamed, Manolya Kavakli-Thorne, and Gulshan Kumar. 2019. "Applications of Generative Adversarial Networks (GANs): An Updated Review." *Archives of Computational Methods in Engineering*, December. `https://doi.org/10.1007/s11831-019-09388-y`.

Banerjee, Rajdeep H., Anoop Rajagopal, Nilpa Jha, Arun Patro, and Aruna Rajan. 2019. "Let AI Clothe You: Diversified Fashion Generation." In *Computer Vision—ACCV 2018 Workshops*, edited by Gustavo Carneiro and Shaodi You, 75–87. Cham: Springer International Publishing.

Baur, Christoph, Shadi Albarqouni, and Nassir Navab. 2018. "Generating Highly Realistic Images of Skin Lesions with GANs" September. `https://arxiv.org/abs/1809.01410`.

Burkov, Andriy. 2019. *The Hundred-Page Machine Learning Book*. Self-published, Amazon.

Cai, Haoye, Chunyan Bai, Yu-Wing Tai, and Chi-Keung Tang. 2018. "Deep Video Generation, Prediction and Completion of Human Action Sequences." In *Computer Vision — ECCV 2018*, edited by Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, 374–90. Lecture Notes in Computer Science. Cham: Springer International Publishing. `https://doi.org/10.1007/978-3-030-01216-8_23`.

Chowdhury, Sohini Roy et al. 2019. "Automated Augmentation with Reinforcement Learning and GANs for Robust Identification of Traffic Signs Using Front Camera Images." In *53rd Asilomar Conference on Signals, Systems & Computers*, 79–83. N.p.: IEEE. `https://doi.org/10.1109/IEEECONF44664.2019.9049005`.

Cohen, Libby. 2020. "Reddit Bans Deepfakes with 'Malicious' Intent." *The Daily Dot*. January 10, 2020. `https://www.dailydot.com/layer8/reddit-deepfakes-ban/`.

Cohn, Gabe. 2018. "AI Art at Christie's Sells for $432,500." *The New York Times*, October 25, 2018. `https://www.nytimes.com/2018/10/25/arts/design/ai-art-sold-c hristies.html`.

Fabbri, Matteo, Simone Calderara, and Rita Cucchiara. 2017. "Generative Adversarial Models for People Attribute Recognition in Surveillance." In *14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. N.p.: IEEE. `https://doi.or g/10.1109/AVSS.2017.8078521`.

Fulgeri, Federico, Matteo Fabbri, Stefano Alletto, Simone Calderara, and Rita Cucchiara. 2019. "Can Adversarial Networks Hallucinate Occluded People With a Plausible Aspect?" *Computer Vision and Image Understanding* 182 (May): 71–80.

Giardina, Carolyn. 2019. "Will Smith, Robert De Niro and the Rise of the All-Digital Actor." *The Hollywood Reporter*, August 10, 2019. `https://www.hollywoodreporter.com/ behind-screen/rise-all-digital-actor-1229783`.

Giles, Martin. 2018. "The GANfather: The Man Who's given Machines the Gift of Imagination." *MIT Technology Review* 121, no. 2 (March/April): 48–53.

Gómez-Bombarelli, Rafael, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. 2018. "Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules." *ACS Central Science* 4, no. 2 (February): 268–76. `https://doi.org/10.1021/acscentsci.7b00572`.

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. "Generative Adversarial Nets." In *Advances in Neural Information Processing Systems*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, 27:2672–2680. Curran Associates, Inc. `https://pr oceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1a fccf3-Paper.pdf`.

Harper, Charlie. 2018. "Machine Learning and the Library or: How I Learned to Stop Worrying and Love My Robot Overlords." *Code4Lib Journal*, no. 41 (August). `https://journa l.code4lib.org/articles/13671`.

Kahng, Minsuk, Nikhil Thorat, Duen Horng Polo Chau, Fernanda B. Viegas, and Martin Wattenberg. 2019. "GAN Lab: Understanding Complex Deep Generative Models Using Interactive Visual Experimentation." *IEEE Transactions on Visualization and Computer Graphics* 25, no. 1 (January 2019): 310–320. `https://doi.org/10.1109/tvcg.2018.28 64500`.

Karras, Tero, Samuli Laine, and Timo Aila. 2019. "A Style-Based Generator Architecture for Generative Adversarial Networks." In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4396–4405. N.p.: IEEE. `https://doi.org/10.1109/ CVPR.2019.00453`.

Korkinof, Dimitrios, Tobias Rijken, Michael O'Neill, Joseph Yearsley, Hugh Harvey, and Ben Glocker. 2018. "High-Resolution Mammogram Synthesis Using Progressive Generative Adversarial Networks." Preprint, submitted July 9, 2018. `https://arxiv.org/abs/ 1807.03401`.

Langr, Jakub and Vladimir Bok. 2019. *GANs in Action: Deep Learning with Generative Adversarial Networks*. Shelter Island, NY: Manning Publications.

Ma, Dongao, Ping Tang, and Lijun Zhao. 2019. "SiftingGAN: Generating and Sifting Labeled Samples to Improve the Remote Sensing Image Scene Classification Baseline In Vitro."

IEEE Geoscience and Remote Sensing Letters 16, no. 7 (July): 1046–1050. `https://doi.org/10.1109/lgrs.2018.2890413`.

Ma, Lei, Yu Liu, Xueliang Zhang, Yuanxin Ye, Gaofei Yin, and Brian Alan Johnson. 2019. "Deep Learning in Remote Sensing Applications: A Meta-Analysis and Review." *ISPRS Journal of Photogrammetry and Remote Sensing* 152 (June): 166–77. `https://doi.org/10.1016/j.isprsjprs.2019.04.015`.

Mazzone, Marian, and Ahmed Elgammal. 2019. "Art, Creativity, and the Potential of Artificial Intelligence." *Arts* 8, no. 1 (March): 1–9. `https://doi.org/10.3390/arts8010026`.

McCormack, Jon, Toby Gifford, and Patrick Hutchings. 2019. "Autonomy, Authenticity, Authorship and Intention in Computer Generated Art." In *Computational Intelligence in Music, Sound, Art and Design*, edited by Anikó Ekárt, Antonios Liapis, and María Luz Castro Pena, 35–50. Cham: Springer International Publishing.

Mukherjee, Sumit, Yixi Xu, Anusua Trivedi, and Juan Lavista Ferres. 2019. "Protecting GANs against Privacy Attacks by Preventing Overfitting." Preprint, submitted December 31, 2019. `https://arxiv.org/abs/2001.00071v1`.

Murphy, Kevin P. 2012. *Machine Learning : A Probabilistic Perspective.* Adaptive Computation and Machine Learning Series. Cambridge, Mass: MIT Press.

Park, Taesung, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. 2019. "Semantic Image Synthesis with Spatially-Adaptive Normalization." In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2332–2341. N.p.: IEEE. `https://doi.org/10.1109/CVPR.2019.00244`.

Radhakrishnan, Sreedhar, Varun Bharadwaj, Varun Manjunath, and Ramamoorthy Srinath. 2018. "Creative Intelligence – Automating Car Design Studio with Generative Adversarial Networks (GAN)." In *Machine Learning and Knowledge Extraction*, edited by Andreas Holzinger, Peter Kieseberg, A Min Tjoa, and Edgar Weippl, 160–75. Cham: Springer International Publishing.

Romera, Eduardo, Luis M. Bergasa, Kailun Yang, Jose M. Alvarez, and Rafael Barea. 2019. "Bridging the Day and Night Domain Gap for Semantic Segmentation." In *2019 IEEE Intelligent Vehicles Symposium (IV)*, 1312–18. N.p.: IEEE. `https://doi.org/10.1109/IVS.2019.8813888`.

Romm, Tony, Drew Harwell, and Isaac Stanley-Becker. 2020. "Facebook Bans Deepfakes, but New Policy May Not Cover Controversial Pelosi Video." *The Washington Post*. January 7, 2020. `https://www.washingtonpost.com/technology/2020/01/06/facebook-ban-deepfakes-sources-say-new-policy-may-not-cover-controversial-pelosi-video/`.

Rostamzadeh, Negar, Seyedarian Hosseini, Thomas Boquet, Wojciech Stokowiec, Ying Zhang, Christian Jauvin, and Chris Pal. 2018. "Fashion-Gen: The Generative Fashion Dataset and Challenge." Preprint, submitted June 21, 2018. `https://arxiv.org/abs/1806.08317`.

Sanchez-Lengeling, Benjamin, and Alán Aspuru-Guzik. 2018. "Inverse Molecular Design Using Machine Learning: Generative Models for Matter Engineering." *Science* 361, no. 6400 (July): 360–365. `https://doi.org/10.1126/science.aat2663`.

Shorten, Connor, and Taghi M. Khoshgoftaar. 2019. "A Survey on Image Data Augmentation for Deep Learning." *Journal of Big Data* 6 (60): 1–48. `https://doi.org/10.1186/s40537-019-0197-0`.

Tirupattur, Praveen, Yogesh Singh Rawat, Concetto Spampinato, and Mubarak Shah. 2018. "Thoughtviz: Visualizing Human Thoughts Using Generative Adversarial Network." In *Proceedings of the 26th ACM International Conference on Multimedia*, 950–958. New York: Association for Computing Machinery. `https://doi.org/10.1145/3240508.3240641`.

Tolosana, Ruben, Ruben Vera-Rodriguez, Julian Fierrez, Aythami Morales, and Javier Ortega-Garcia. 2020. "DeepFakes and Beyond: A Survey of Face Manipulation and Fake Detection." Preprint, submitted January 1, 2020. `https://arxiv.org/abs/2001.00179`.

Uřičář, Michal, Pavel Křížek, David Hurych, Ibrahim Sobh, Senthil Yogamani, and Patrick Denny. 2019. "Yes, We GAN: Applying Adversarial Techniques for Autonomous Driving." In *IS&T International Symposium on Electronic Imaging*, 1–16. Springfield, VA: Society for Imaging Science and Technology. `https://doi.org/10.2352/ISSN.2470-1173.2019.15.AVM-048`.

Vondrick, Carl, and Antonio Torralba. 2017. "Generating the Future with Adversarial Transformers." In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2992–3000. N.p.: IEEE. `https://doi.org/10.1109/CVPR.2017.319`.

Wang, Zhengwei, Qi She, and Tomas E. Ward. 2019. "Generative Adversarial Networks: A Survey and Taxonomy." Preprint, submitted June 4, 2019. `https://arxiv.org/abs/1906.01529`.

Wen, Shiping, Weiwei Liu, Yin Yang, Tingwen Huang, and Zhigang Zeng. 2019. "Generating Realistic Videos From Keyframes With Concatenated GANs." *IEEE Transactions on Circuits and Systems for Video Technology* 29 (8): 2337–48. `https://doi.org/10.1109/TCSVT.2018.2867934`.

Yi, Xin, Ekta Walia, and Paul Babyn. 2019. "Generative Adversarial Network in Medical Imaging: A Review." *Medical Image Analysis* 58 (December): 1–20. `https://doi.org/10.1016/j.media.2019.101552`.

Zhavoronkov, Alex. 2018. "Artificial Intelligence for Drug Discovery, Biomarker Development, and Generation of Novel Chemistry." *Molecular Pharmaceutics* 15, no. 10 (October): 4311–13. `https://doi.org/10.1021/acs.molpharmaceut.8b00930`.

Zhu, Jun-Yan, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks." In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2242–2251. N.p.: IEEE. `https://doi.org/10.1109/ICCV.2017.244`.