

Using GANs to Synthesise Minimum Training Data for Deepfake Generation

Simranjeet Singh¹, Rajneesh Sharma¹, and Alan F. Smeaton^{1,2}

¹School of Computing and ²Insight Centre for Data Analytics
Dublin City University, Glasnevin, Dublin 9, Ireland
alan.smeaton@dcu.ie

Abstract. There are many applications of Generative Adversarial Networks (GANs) in fields like computer vision, natural language processing, speech synthesis, and more. Undoubtedly the most notable results have been in the area of image synthesis and in particular in the generation of deepfake videos. While deepfakes have received much negative media coverage, they can be a useful technology in applications like entertainment, customer relations, or even assistive care. One problem with generating deepfakes is the requirement for a lot of image training data of the subject which is not an issue if the subject is a celebrity for whom many images already exist. If there are only a small number of training images then the quality of the deepfake will be poor. Some media reports have indicated that a good deepfake can be produced with as few as 500 images but in practice, quality deepfakes require many thousands of images, one of the reasons why deepfakes of celebrities and politicians have become so popular. In this study, we exploit the property of a GAN to produce images of an individual with variable facial expressions which we then use to generate a deepfake. We observe that with such variability in facial expressions of synthetic GAN-generated training images and a reduced quantity of them, we can produce a near-realistic deepfake videos.

Keywords: Deepfake generation; Generative Adversarial Networks; GANs; Variable face images.

1 Introduction

Recently we have seen a rise in the presence of deepfake videos on social media and in entertainment applications. Sometimes these are used for good but it is the mis-use of deepfakes which attracts most media attention and commentary. What makes deepfakes so important today is their low barrier to entry, meaning that easily available tools and models can be used by researchers with even moderate programming skills to generate very realistic deepfake videos. When this is considered in the context of targeted advertisements for political elections on social media, then the impact of deepfakes could be quite significant.

A deepfake is a video created by manipulating an original video using advanced machine learning techniques. This involves replacing the face of an individual from a source video with the face of a second person in the destination video. A model of

the face of the second person, the one who is superimposed into the destination video, is created based on a typically large collection of facial images. In the early days of deepfake videos, celebrities were used in the destination videos because (a) it is easy to get thousands of images of celebrities from the internet and (b) most of these pictures are of the subject facing the camera. The Hollywood actor Nicholas Cage became even more of a celebrity as a model based on images of his face was one of the first to be made publicly available and was widely used in creating deepfakes when the interest was in the quality of the generated videos and less on who the subjects were.

Now that we have reached the point where the quality of deepfakes is almost indiscernible from real videos, interest returns to how to generate these deepfakes, not using celebrities as the subjects but using ordinary people. While there are nefarious applications based on the use of deepfakes of non-celebrity individuals, there are also useful scenarios. An example of this is using deepfake videos of a non-celebrity as a sales agent or troubleshooter in an online chat system.

One characteristic of the non-celebrity subject in a deepfake, is that there will typically be a limited number of images of the subject's face available for training a deepfake generator, perhaps even no images to start from. Thus we expect that training data, i.e. images of the face, may actually be taken from short video clips recorded specifically for this purpose.

In this paper we look at how deepfake videos of non-celebrity subjects can be generated using limited training data, i.e. a small number of training images. In particular we are interested not just in the limited number of images used but also in the variability of facial expressions among those limited number of images. To test this we use a large number of images to create a model of an individual face, and then we generate a small number of synthetic but realistic images from that model which we use to generate a deepfake. While this may seem counter intuitive, to use a large number of images of a celebrity to generate a small number of synthetic images of that celebrity this allows the synthetic images to include a lot of facial variety of expression which we could not obtain easily if we were to use a real collection as the small number of deepfake training images.

The rest of this paper is organised as follows. In the next section we present an overview of Generative Adversarial Networks (GANs) followed by a description of 4 metrics used to evaluate the quality of our output from image-generating GANs. We then describe how we gathered, or more correctly how we synthesised image data for training a GAN and we then present an analysis of those images in terms of their quality and variability of facial expressions. That is followed by a description of how we used those images to create a deepfake and then some conclusions and plans for future work.

2 Generative Adversarial Networks (GANs)

The idea behind adversarial networks was first published by Olli Niemitalo however his ideas were never implemented [24] and a similar concept was introduced by Li, Gauci and Gross in 2013 [24]. Generative Adversarial Network (GAN) implementations were first described in 2014 by Ian Goodfellow and until 2017 the use of GANs was restricted to just image enhancement to produce high quality images. In 2017 GANs were used for

the first time for generating new facial images and the idea began to make its presence known in the fine arts arena and were thus dubbed creative adversarial networks [24].

GANs have been widely applied to domains such as computer vision, natural language processing, etc. GANs have contributed immensely to the field of image generation [16] where the quality of synthetic images a GAN can produce has improved significantly over the years since its inception. Other example applications of GANs include the generation of DNA sequences, 3D models of replacement teeth, impressionist paintings, and of course video clips, some known as deepfakes.

Deepfakes are a form of video manipulation where two trained networks are pitted against each other to generate an output of sufficient quality as to be close to indecipherable. They operate by inputting a set of images of a subject from which they build a model of the face and then superimpose this face model on the target face in an original video.

One of the challenges faced by deepfake generation, apart from their computational cost, is the requirement for a large number of training images of the subject to be faked into the original image. In practice, the quality of the generated deepfake will depend not only on the number of face images in the training data but the amount of facial variability among those images and the amount of facial variation in the original video. If the original video has a face with not much emotion shown and very little variation in facial expression then it follows that the training data for the face to be superimposed does not need a wide variety of facial expression and thus a smaller number of training images are needed. If the original video has a lot of facial variation then the model to be generated to replace this original face will need to be larger and more complex, and thus require far more training data. Some commentators have said that as few as 500 images of the face of a subject are required for a good deepfake but in practice these refer to deepfakes without much facial emotion and the best deepfakes are generated using many thousands of source images of the face.

Deepfakes have many applications in the entertainment industry such as movie production and the Salvador Dali museum in St Petersburg, Florida¹, but there are also applications in areas like customer relations where text or audio chatbots are replaced by real people or deepfakes, or in assistive technology where older people living alone might interact with generated media which could consist of deepfaked videos of loved ones. The problem with such applications is that there are usually few images available from which to train a model to create a deepfake.

In this study we look into how the amount of, and the variety of facial expressions included in, the face data used to train a deepfake generator affects the quality of the deepfake. One of the latest GANs, StyleGAN2 [16], is used in our study to produce synthetic facial images for training and various evaluation methods are used to benchmark the quality of these synthetic images including Inception score [21] and the Fréchet Inception Distance [14], and the variety among those faces using OpenFace’s Comparison method [2] and *face_recognition*’s compare method [13]. Our implementation of StyleGAN2 is trained on a dataset of 132,000 images taken from stills of YouTube videos of TV night show host John Oliver and from this we synthesise 1,000 images in a way that includes a lot of facial variation. We then use these 1,000 images to train deepfake-

¹ <https://thedali.org/>

lab [19] to generate a deepfake where the (synthesised) John Oliver is superimposed on a subject shown interacting with a chatbot in a dialogue. A schematic of the flow of our data processing is shown in Figure 1.

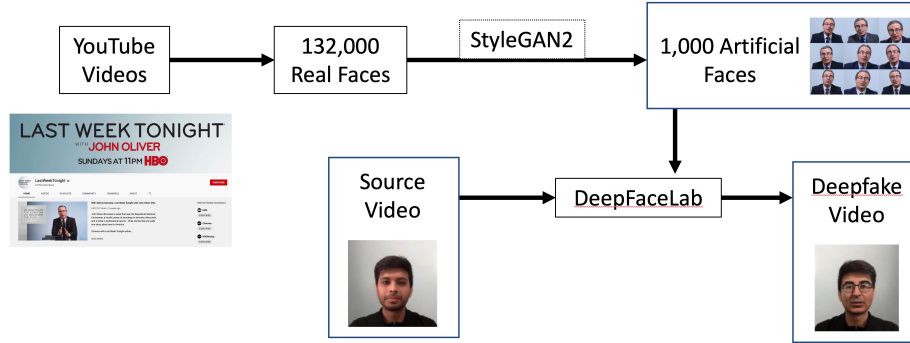


Fig. 1. Overview of data processing in the paper

As we show later in this paper, when trained with enough facial variations in input images, we found that deepfacelab is able to produce an accepted quality of generated deepfakes.

3 Evaluation Metrics

There are a number of methods developed to evaluate the quality of output produced by GANs and to measure the variability in a set of images of faces and we discuss some of these here. For a more detailed description of GAN output see [4].

3.1 Inception Score (IS)

Inception Score was first introduced by Salimans *et al.* [21], and is the most common method used for evaluating GAN outputs. It uses a pre-trained inception model to classify generated images and calculates probabilities of each image belonging to each class, and looks at the label distribution. Images with high probability towards one class/label are considered high quality.

In summary, Inception Score actually captures two properties of a generated dataset:

1. Image Quality: How highly an image belongs to one class as classified by an inception classifier . . . do they look similar to a specific object?
2. Image Diversity: How many different images are generated by the GAN . . . is there a range of different objects generated?

Inception score has a lowest value of 1.0 and higher values indicate an improving quality of the GAN [7]. However, even with these properties, IS has its limitations

as shown in [4]. Firstly, it favours GANs which can store training data and generate images around centers of data modes and secondly, since this method uses an Inception Classifier which is trained on the ImageNet dataset with many object classes, it may uplift those models which produce good images of objects. A third limitation of IS is that since the score never takes a real dataset into account and evaluates the quality of a GAN based on its generated dataset, this can be deceptive. This may favour GANs which produce clear and diverse images of any object, far from a real dataset.

3.2 Fréchet Inception Distance (FID)

FID is another popular method for GAN evaluation introduced by Heusel *et al.* in 2017 [14]. It uses feature vectors of real data and generated data and calculates distances between them. The FID score is used to evaluate the quality of images generated by GANs, and lower scores have been shown to correlate well with higher quality generated images [6].

Unlike Inception Score (IS), FID captures the statistics of generated data and compares it with the statistics of real data. It is similar to IS in the way that it also uses the inception v3 model. Instead of using the last output layer of the model, it uses the last coding layer to capture specific features of the input data. These are collected for both real and generated data. The distance between two distributions, real and generated, is then calculated using Fréchet-distance [1] which itself uses the Wasserstein-2 distance which is a calculation between multi-variate Gaussians fitted to data embedded into a feature space [4]. Lower distance values convey that the generated dataset is of high quality and similar to real dataset [6].

A model that generates only one image per class will have a bad FID score whereas the same case will have high IS. FID compares data between real and generated data sets whereas IS only measures diversity and quality of a generated dataset. Unlike IS, data scores will be bad on an FID scale in cases where there is noise or other additions to the data [4].

3.3 OpenFace Python Library

OpenFace is an open source general-purpose library for face recognition [2] with various features including dlib's face landmark detector [17]. Landmarks are used to crop images to ensure only facial data is passed to the neural network for training, producing a low-dimensional face representation for the faces in images [2]. OpenFace includes a function to calculate the squared L2 distance [5] between facial representations, providing a comparison function among faces in a dataset. An image in the dataset can be paired with every other image in the dataset and the squared L2 distance computed, ranging from 0 to 4, with 0 meaning the faces in two compared images are more likely to be of the same person [3].

In our work we applied OpenFace to an image set generated by StyleGAN2 to measure the degree of variability among the generated faces and we computed the mean and variance of inter-image scores among the images. To confirm our approach, two datasets of facial images were generated, each with 100 images of the same person taken from a smartphone in burst mode. In one dataset, the facial expressions were

kept the same and we called this dataset the “Monotone” dataset. In the second dataset, various facial expressions were captured called the “Varied” dataset. The number of comparisons this requires is 4,950 for each dataset from which we compute mean and variance.

Figure 2 shows a subset of each dataset with calculated mean and variance in Table 1. The Monotone dataset gave a smaller mean and variance score which denotes the person in the dataset is same but with less variation in facial expression compared to the other dataset which has variability in facial expressions of the individual, though since the mean is still close to zero, the person in the dataset is the same person.



Fig. 2. Subset of Monotone and Varied facial image datasets

Dataset	Mean	Variance
Monotone faces	0.019	0.000
Varied faces	0.234	0.019

Table 1. Mean and variance among two 100 facial image datasets

3.4 The face_recognition Python Library

face_recognition is a simple library in Python for face recognition which also uses dlib’s facial landmark detector [17] and has a comparison feature which calculates distance between facial landmarks of two images. Given a certain threshold, it returns a True/False whether the person in both images is the same or not [13]. To show its capabilities power, we compared two images of the same individual shown in Figure 3, the first taken in 2007 and the second in 2019 and [13] detects these as the same person. For our purposes we iterate through the GAN-generated images and compare each with the original images used to train StyleGAN2 using *face_recognition* as another way of evaluating the GAN-generated dataset.



Fig. 3. Two images of the same individual taken 12 years apart.



Fig. 4. 10 images of some celebrities taken years apart. The first row are images from earlier in their careers and the second row are more recent images.

To further validate this method, we took 10 pairs of celebrity face images shown in Figure 4, each pair of images taken years apart [22] and using *face_recognition* we compared them, observing that each pair is identified by *face_recognition* as of same person.

4 Facial Image Data Gathering for GAN Training

For training a deepfake video generation system there are numerous datasets available from other studies [11] where facial data was gathered but for almost all of these the data was either, not in sufficient quantity or quality to train a GAN, or consists of faces of different individuals whereas we require images of the same person.

The GAN we use is StyleGAN2 developed by Karras *et al.* in 2019 with improvements over its predecessor StyleGAN [16] [15]. StyleGAN2 can generate images up to 1024x1024 pixels in size but this requires hardware intensive training. We worked at 256x256 pixels image resolution considering the limited hardware available for this study and we generated our own dataset by extracting frames from videos of an individual.

As stated in [23], John Oliver is “an English-American comedian, writer, producer, political commentator, actor, and television host”. He is the host of a popular HBO Series “Last week tonight with John Oliver” [23]. We chose to use videos of him because he is always in the frame and at same position on screen and talks with various facial expressions. His recent videos have a plain background because of being shot at a home studio due to COVID-19.

Using 20 videos from the official YouTube channel² we extracted 132,000 frames cropped to John Oliver’s face area with the remaining part of the frames ignored. We resized the images to 256x256 pixels for model training, using the Pillow python library [20]. We trained StyleGAN2 [16] by converting to TFRecords format [12] which took around 30 minutes of processing and around 27GB of storage on a system with 30GB of memory and 1 NVIDIA Tesla V100 GPU on the Google Cloud Platform.

StyleGAN2 training uses the TFRecord format [12] for storing a sequence of binary records which is advantageous in cases of large datasets which cannot be held in memory during training and only require the dataset one at a time (e.g. a batch) when they loaded from disk and processed [18]. The original dataset used by StyleGAN2 training had 70K images from the Flickr Faces HQ (FFHQ) [16] at 1024x1024 repeated 25M times. That means if the dataset contains 70K images, the GAN would go over these repetitively $25M/70K \simeq 357$ to learn the salient features from 25M images. The authors state they performed training with 8 GPUs for almost 10 days to generate high quality images [18].

Using our John Oliver dataset of 132,000 images, training was completed with the number of images set to 500K. This made the GAN go over the dataset only $500K/132K \simeq 3.8$ times. Since the dataset size was significant and had variation in images even though all images are of a single person, the GAN was able to generate quality output images and Figure 5 shows some of these images.



Fig. 5. Sample of generated images from John Oliver’s trained model

After generating 1,000 images of John Oliver from our trained model, we applied the metrics described in Section 3 and the results are shown in Table 2. Out of 1,000 generated faces, the face_recognition library recognised 929 faces as John’s faces, re-

² <https://www.youtube.com/user/LastWeekTonight>

jected 66 images and 5 of the 1,000 images were not recognised as faces due to having noise in them around the facial landmarks.

Metric	Value
Inception Score \uparrow	Mean: 1.635 Variance: 0.011
Fréchet Inception Distance \downarrow	29.318
OpenFace library comparison \downarrow	Mean: 0.401 Variance: 0.068
face_recognition library comparison	Matching: 929 Not matching: 66 Not faces: 5

Table 2. Scores for 1,000 image output from StyleGAN2 trained on John Oliver dataset, \uparrow means higher is better and \downarrow means lower is better.

These figures gives us certainty that the images generated by the GAN are recognised as John Oliver’s face. The comparison feature of the OpenFace Library [2] gave a mean score of 0.401 and variance of 0.068 which means that the facial data is of the same person and it has sufficient variability. The calculated Inception score [21] for the same dataset was 1.628.

5 Deepfake Video Generation

we used an NLP-based chatbot python library Chatterbot [8] to generate dialogues which we could use for our source videos. We built a simple bot and trained it on an English language corpus which has various topics of conversation [9]. Since this was built to give dialogue for our videos, we generated only a few dialogues with the bot and one conversation with random queries after training goes as follows:

Human: Do you know gossip?

Bot: Gregory said I respond to the current line, not with respect to the entire conversation. Does that count as gossip?

Human: What is AI?

Bot: Artificial Intelligence is the branch of engineering and science devoted to constructing machines that think.

Human: What is a super computer?

Bot: Computers which can perform very large numbers of calculations at very high speed and accuracy are called super computers.

Human: How does a computer work?

Bot: Computers are very dumb. They only execute instructions given by humans.

Human: What was the first computer?

Bot: It's hard to say, but The ENIAC is regarded as the first 'real' computer. It was developed at University of Pennsylvania in 1946.

We recorded a video with our own subject responding in the above dialogue. The DeepFaceLab [19] face swapping framework was then used to swap the face of our subject with faces from the 1,000 GAN-generated facial images of John Oliver. DeepFaceLab's model was trained to find the facial area of our subject in our recorded video while the GAN-generated John Oliver facial images were fed to the model. The model was trained for almost 80,000 iterations to learn the features from our subject's face and John Oliver's face and swap them. The output of this swap was a regenerated video dialogue at 256x256 pixel resolution with John Oliver's face having same dialogue as our subject, in other words it is a deepfake. A still image is shown in Figure 6 and the video is available for anonymous viewing at <https://bit.ly/31xEjgy>



Fig. 6. Still frame from video illustrating face swapping (deepfake) from our subject (left side) to John Oliver (right side), video is available at <https://bit.ly/31xEjgy>

6 Conclusions and Future Work

In this study we introduced and developed an idea to consolidate various techniques available to develop a video dialogue of an individual speaking to camera based on a limited dataset of images of the individual. We generated a dataset of 132,000 video frames extracted from TV night show host John Oliver's YouTube videos and trained the StyleGAN2 [16] GAN to generate a sample of 1,000 images and 4 evaluation methods were used to measure the variability and quality of these images. These included the Python libraries OpenFace and face_recognition which measure facial variability in a dataset of faces.

We then generated several dialogues from a chatbot we trained and recorded a video with our own subject responding as part of one of these dialogues. We applied a Face Swapping Framework DeepFaceLab [19] to swap the face of our subject with that of the GAN-generated John Oliver images. The final video output of swapped dialogues alongside the original dialogues is publicly and anonymously available at <https://bit.ly/31xEjgy>.

We observe that the deepfake video based on a synthetic set of 1,000 images of John Oliver is of good quality. There is some colour variation across frames which we could easily have smoothed using a tool like OpenCV but we decided to leave it there to emphasise to the viewer how the video was created.

Our future work is to repeat the video generation process using a more homogeneous set of images generated by the GAN which synthesises images of John Oliver, and then to compare the quality of the generated deepfakes. While most work on deepfakes has been to detect them, such as [10], there is little work reported to date on measuring their quality so ultimately the measure of deepfake quality may be how easily it is to be recognised as a deepfake.

Acknowledgments. We wish to thank Satyam Ramawat for acting as a test subject for our image generation and AS is part-funded by Science Foundation Ireland under grant number SFI/12/RC/2289.P2, co-funded by the European Regional Development Fund.

References

1. Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications*, 5(01n02):75–91, 1995.
2. Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. Technical report, CMU-CS-16-118, CMU School of Computer Science, 2016.
3. Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. *OpenFace: Free and open source face recognition with deep neural networks*, 2016 (accessed July 07, 2020). <https://cmusatyalab.github.io/openface/demo-2-comparison/>.
4. Ali Borji. Pros and cons of GAN evaluation measures. *Computer Vision and Image Understanding*, 179:41–65, 2019.
5. Jason Brownlee. *Gentle Introduction to Vector Norms in Machine Learning*, 2018 (Last accessed July 26, 2020). <https://machinelearningmastery.com/vector-norms-machine-learning/>.
6. Jason Brownlee. *How to Implement the Fréchet Inception Distance (FID) for Evaluating GANs*, 2019 (Last accessed June 26, 2020). <https://machinelearningmastery.com/how-to-implement-the-frechet-inception-distance-fid-from-scratch>.
7. Jason Brownlee. *How to Implement the Inception Score (IS) for Evaluating GANs*, 2019 (Last accessed June 26, 2020). <https://machinelearningmastery.com/how-to-implement-the-inception-score-from-scratch-for-evaluating-generated-images>.
8. Gunther Cox. Building chatbot using chatterbot, (Last accessed Aug 09, 2020). <https://chatterbot.readthedocs.io/en/stable/>.
9. Gunther Cox. *ChatterBot Language Training Corpus*, (Last accessed Aug 09, 2020). https://github.com/gunthercox/chatterbot-corpus/tree/master/chatterbot_corpus/data/english.
10. Brian Dolhansky, Russ Howes, Ben Pflaum, Nicole Baram, and Cristian Canton Ferrer. The deepfake detection challenge (DFDC) preview dataset. *arXiv preprint arXiv:1910.08854*, 2019.

11. Bob Fisher. *CVonline: Image Databases*, 2019 (Last accessed July 26, 2020). <http://homepages.inf.ed.ac.uk/rbf/CVonline/Imagedbase.htm#face>.
12. Thomas Gamauf. *Tensorflow Records? What they are and how to use them*, 2018 (Last accessed July 28, 2020). <https://medium.com/mostly-ai/tensorflow-records-what-they-are-and-how-to-use-them-c46bc4bbb564>.
13. Adam Geitgey. *face_recognition*, 2017 (accessed July 07, 2020). https://github.com/ageitgey/face_recognition/.
14. Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017.
15. Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
16. Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020.
17. Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009.
18. NVidia Labs. *StyleGAN Github*, 2019 (Last accessed July 26, 2020). <https://github.com/NVlabs/stylegan2>.
19. Ivan Petrov, Daiheng Gao, Nikolay Chervoniy, Kunlin Liu, Sugasa Marangonda, Chris Umé, Jian Jiang, Luis RP, Sheng Zhang, Pingyu Wu, et al. DeepFaceLab: A simple, flexible and extensible face swapping framework. *arXiv preprint arXiv:2005.05535*, 2020.
20. python pillow. *Pillow, friendly PIL fork*, (Last accessed July 26, 2020). <https://github.com/python-pillow/Pillow>.
21. Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
22. S. Sengupta, J.C. Cheng, C.D. Castillo, V.M. Patel, R. Chellappa, and D.W. Jacobs. Frontal to profile face verification in the wild. In *IEEE Conference on Applications of Computer Vision*, February 2016.
23. Wikipedia. *John Oliver*, Last accessed July 28, 2020. https://en.wikipedia.org/wiki/John_Oliver.
24. Wikipedia. *Generative adversarial network*, (Last accessed on May 03, 2020). https://en.wikipedia.org/wiki/Generative_adversarial_network#History.