Available online www.ejaet.com

European Journal of Advances in Engineering and Technology, 2019, 6(11):1-9

Research Article ISSN: 2394 - 658X

# Large Scale High Fidelity Facial Recognition Security Architecture with Personalized GAN Trained Discriminator

**Mamoun Abu-Samaha and Jeong Hee Kim**

*International Technological University, 2711 N 1st St., San Jose, CA 95134*

**ABSTRACT**

Decisions based on machine learning algorithms potentially impact billions of human lives. They must be robust, reliable, accurate, usable, and understandable. This research investigated key theoretical pillars of machine learning's generalization, training, and test accuracy. The focus was the extent to which empirical behavior matches existing theory. The work produced novel methods for enhancing accuracy and generalization. Reproducible, observed behaviour was characterized across differences in optimization algorithm, dataset, task, evaluation metric, and architecture.

Optimization algorithms bias machine-learning models toward solutions with varying generalization properties. A transpose CNN in the generator adaptive algorithm empirically found solutions with enhanced features, with minimal data and better behavior, compared to using a stochastic gradient descent and leveraging GAN. The GAN technique creates mode collapse issues and requires more samples, which ultimately yields less accurate results.

We constructed an example using a simple convolutional neural network in the generator, part of the GAN model corroborating the algorithm's empirical behavior on neural networks. We studied machine learning models' generated images from commonly used datasets in academic benchmarks and in machine learning competitions.

New test sets for real images as datasets allowed us to evaluate a broad range of classification models without mode collapse issues, but with improved accuracy. Previously generated accuracy numbers were susceptible to even minute natural variations in the data distribution.

A layered architecture of discriminators; one general CNN discriminator and one personalized per object GAN, trained on the required feature(s), was optional. It increased accuracy and improved recognition features. The architecture and method across different data domains, loss functions, model classes, and human analysts was robust.

Robust machine learning in facial recognition problems should focus on designing distributed multilayer discriminators and advanced GAN-based transpose convolution in the generator. Creating a model to work in dynamic environments is a challenging, but necessary, undertaking.

**Key words:** generative adversarial networks, facial recognition, transposed convolutional neural networks, multilayered discriminators, multilayer perceptrons, deep learning

## 1. INTRODUCTION

In our presentation here, we first examine how optimization algorithms bias machine-learning models towards solutions with varying generalization properties. We show that the use of a transpose CNN in the generator adaptive algorithm empirically finds solutions with enhanced features with minimal data and better behaviour when compared to those found by stochastic gradient decent and leveraging GAN. The use of a GAN technique creates mode collapse issues and requires more samples, which ultimately yields less accurate results.

Next, we construct an example using a simple transverse convolutional neural network in the generator, which is part of the GAN model that corroborates the algorithm's empirical behaviour on neural networks.

Following this step, we then study the extent to which machine learning models have generated images from both commonly used datasets in both academic benchmarks and machine learning competitions.

In this process, we build new test sets for real images as datasets and evaluate a broad range of classification models on the new datasets. Thus far, we have had no collapse mode issues when we test these datasets and we also have improved accuracy.

In studying image datasets in general, we have seen that previously generated accuracy numbers are susceptible to even minute natural variations in the data distribution.

Surprisingly, despite several years of adaptively selecting the models to perform well on these competitive benchmarks, we find no evidence of collapse mode. In our solution, we find that a layered architecture of discriminators; one general CNN discriminator and one personalized per object GAN, trained on the required feature(s) was optional and both increased accuracy and improved the features of recognition. These findings speak to the robustness of the architecture and the method across different data domains, loss functions, model classes, and human analysts.

Overall, our research suggests that the true concern for robust machine learning in facial recognition problems should be focused on the design of distributed multilayer discriminators and advanced GAN-based transpose convolution in the generator. In concluding our work on the design of a new model for facial recognition and detection, we find that to create a model which will reliably work in dynamic environments is a challenging, but necessary undertaking.

**Formal background on GAN and transpose CNN**

In this section, I present background information regarding the basic concept of "generative adversarial networks" (GANs) and that of "transpose convolutional neural networks (CNNs).

Since their introduction in 2014, as part of the work of Ian Good fellow and his Ph.D. advisors, Generative Adversarial Networks (GANs) have gained stature as one of the most widely used methods for training deep generative models. GANs have been successfully deployed for tasks such as photo super-resolution, object generation, video prediction, language modelling, vocal synthesis, and semi-supervised learning, amongst many others.

At the core of the GAN methodology is the idea of jointly training two networks: a generator network, meant to produce samples from some distribution (that ideally will mimic examples from the data distribution), and a discriminator network, which attempts to differentiate between samples from the data distribution and the ones produced by the generator. This problem is typically written as a min-max optimization problem of the following form: min G max D (Ex∼pdata [log D(x)] + Ez∼platent [log(1 − D(G(z))]).

***Despite the preference for GANs as a tool, the actual task of optimizing GANs remains a challenging problem, both from a theoretical and a practical standpoint.***

Although the original GAN paper included some analysis on the convergence properties of the approach, it assumed that updates occurred in pure function space, allowed arbitrarily powerful generator and discriminator networks, and modeled the resulting optimization objective as a convex-concave game, therefore yielding well-defined global convergence properties.

Furthermore, this analysis assumed that the discriminator network is fully optimized between generator updates, an assumption that does not mirror the practice of GAN optimization. Indeed, in practice, there exist a number of well-documented failure modes for GANs such as mode collapse or vanishing gradient problems.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.arXiv:1706.04156v3, 13 Jan 2018.

https://arxiv.org/pdf/1706.04156.pdf

**GAN Definition**

https://searchenterpriseai.techtarget.com/definition/generative-adversarial-network-GAN

A generative adversarial network (GAN) is a type of AI machine learning (ML) technique made up of two nets that are in competition with one another in a zero-sum game framework. GANs typically run unsupervised, teaching itself how to mimic any given distribution of data.

The two neural networks that make up a GAN are referred to as the generator and the discriminator. The generator is a type of convolutional neural network (footnote) that will create new instances of an object, and the discriminator is a type of deconvolutional neural network (footnote) that will determine its authenticity, or whether or not it belongs in a dataset. Both entities compete during the training process where their losses push against each other to improve behaviors, known as backpropagation (footnote). The goal of the generator is to produce passable output without being caught, while the goal of the discriminator is to identify the fakes. As the double feedback loop continues, the generator produces higher-quality output and the discriminator becomes better at flagging imposters.

**Definition/Description of Transpose CNN**

**Deep Learning and Convolutional Neural Networks**

**Definition of "Deep Learning"**

In deep learning, a subfield of AI research, the computer is not explicitly told how to solve a particular task such as object classification. Instead, it uses a learning algorithm to extract patterns in the data that relate the input data, such as the pixels of an image, to the desired output such as the label "cat." The challenge for researchers has been to develop effective learning algorithms that can modify the weights on the connections in an artificial neural network so that these weights capture the relevant patterns in the data.

**Definition of "Convolutional Neural Network"**

A **Convolutional Neural Network (ConvNet/CNN)** is a Deep Learning algorithm which can take in an input image, assign importance (learnable weights and biases) to various aspects/objects in the image and be able to differentiate one from the other [22].

Furthermore, *"The convolutional neural network (CNN) is a class of* ***deep learning neural networks****. CNNs represent a huge breakthrough in image recognition. They're most commonly used to analyze visual imagery and are frequently working behind the scenes in image classification.. . . Image classification is the process of taking an* ***input*** *(like a picture) and outputting a* ***class*** *(like "cat") or a* ***probability*** *that the input is a particular class ("there's a 90% probability that this input is a cat")."* [

**Wikipedia definition of "deep learning"**
https://en.wikipedia.org/wiki/Deep_learning
**Wikipedia definition of "convolutional neural network"**
https://en.wikipedia.org/wiki/Convolutional_neural_network#Convolutional

Furthermore, in deep learning, a **convolutional neural network** (**CNN**) is a class of deep neural networks, most commonly applied to analyzing visual imagery.

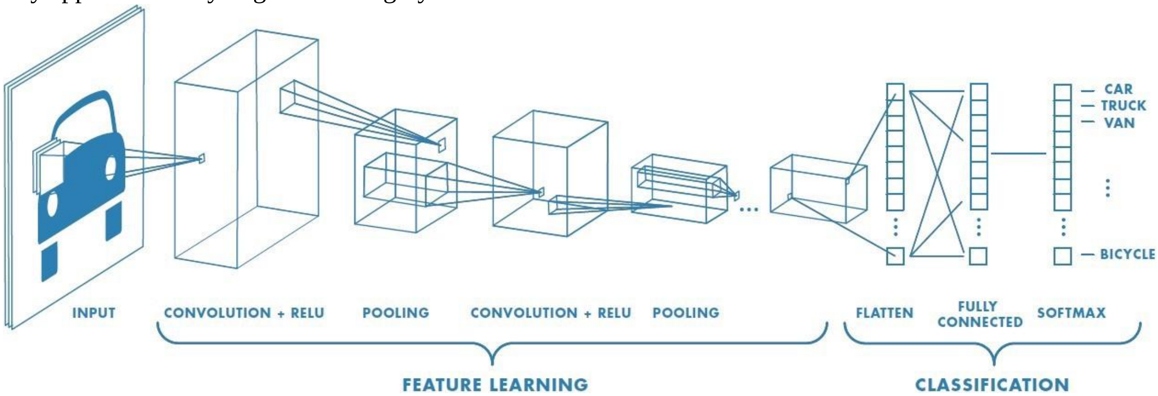


**Fig. 1** Example of a Convolutional Neural Network

A convolutional neural network (CNN, or ConvNet) is one of the most popular algorithms for deep learning with images and video. Like other neural networks, a CNN is composed of an input layer, an output layer, and many hidden layers in between.

**Feature Detection or Learning**

Feature Detection Layers These layers perform one of three types of operations on the data: convolution, pooling, or rectified linear unit (ReLU).

Convolution puts the input images through a set of convolutional filters, each of which activates certain features from the images.

Pooling simplifies the output by performing nonlinear down sampling, reducing the number of parameters that the network needs to learn about.

Rectified linear unit (ReLU) allows for faster and more effective training by mapping negative values to zero and maintaining positive values. These three operations are repeated over tens or hundreds of layers, with each layer learning to detect different features.

**Classification Layers**

After feature detection, the architecture of a CNN shifts to classification.

The next-to-last layer is a fully connected layer (FC) that outputs a vector of K dimensions where K is the number of classes that the network will be able to predict. This vector contains the probabilities for each class of any image being classified.

The final layer of the CNN architecture uses a Softmax function to provide the classification output.

**Transpose Convolutional Neural Network (CNN) – What is It and How Does It Work?**

https://medium.com/apache-mxnet/transposed-convolutions-explained-with-ms-excel-52d13030c7e8

One of the best ways for us to gain some intuition about the nature of transpose CNN is by looking at examples from Computer Vision that use the transposed convolution method. Most of these examples start with a series of regular convolutions to ***compress*** the input data into an abstract spatial representation, and then use transposed convolutions to ***decompress*** the abstract representation into something of use.

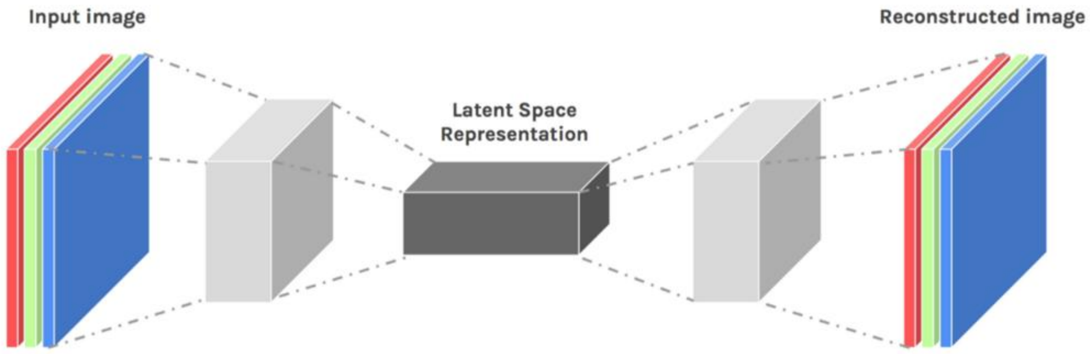


**Fig. 2**

A convolutional auto-encoder is tasked with recreating its input image, after passing intermediate results through a ‘bottleneck’ of a limited size. Uses of auto-encoders include compression, noise removal, colorization and in-painting. Success depends on being able to learn dataset specific compression in the convolution kernels and dataset specific decompression in the transposed convolution kernels.

**Overview**

**Part I.** How optimization algorithms bias machine learning models towards solutions with varying generalization properties.

**Part II.** The Empirical Behavior Match with Theory: Accuracy

A. The use of a transpose CNN produces solutions with enhanced features using minimal data results in better behavior than methods using stochastic gradient descent (SGD).
B. In contrast, the use of a generative adversarial network (GAN) technique creates mode collapse issues and requires more samples.
C. Furthermore, the need for more samples ultimately yields less accurate results.

**Part III.** Confirmation of the new behavioral model using a simple transverse convolution neural network (TCNN) in the generator

A. We obtain successful corroboration that the GAN model is consistent with the algorithm’s empirical behavior on neural networks.

**Part IV.** Discussion regarding the extent to which machine learning models have generated images from both commonly used datasets in both academic benchmarks and machine learning competitions.

A. The construction of new datasets and evaluation of a broad range of classification models on these new datasets.
B. Overall result: No collapse mode issues and improved accuracy.
C. A review of past image datasets shows that previously generated accuracy is susceptible to even minute natural variations in the data distribution.

**Part V.** Conclusions

A. Unrolled optimization of a GAN discriminator does not, contrary to general belief, eliminate instability of the generative adversarial network.
B. Large images optimization is required for images used for security application.
C. Multi-layer and personalized Discriminators will increase accuracy
D. Using GAN, syntactic data, provide large number of training data, covering un-detectable by human error.

**WAN-GAN Architecture**

The probability distribution of data might be complicated and very hard, hence poor results

So, having a generative machine that can generate samples from noise without having to deal with nasty probability distribution is a very compelling reason for using GAN.

With GAN we get sample data relatively cheaply using the trained Generative Adversarial Network.

**Facial Recognition Security Using GAN**

1- Extract key signature features through GAN for each person.
2- Build a discriminator for features found in Step 1.
3- Build a database of all individuals in a group using Steps 1 and 2.
4. Identify a person by taking a video clip of the person’s face rotating positions of camera.
5- Repeat Step 1 for person to be identified using images from Step 4.
6- Classify the face image of the unknown person by search through database built in Step 3.
7- Build a network to provide the captured image data to each group on parallel

**Advantages**

The probability distribution of data might be a very complicated distribution and very hard and intractable to infer. So, having a generative machine that could generate samples from *data* without having to deal with nasty probability distribution is very nice.

If we have this, then we can use it for another process that requires a sample from *Pdata* as we can get samples relatively cheaply using the trained Generative Net.

The architecture is designed to be simple. Since the GAN and discriminator are done for a person rather than a large group of people, it can be simplified, e.g. DNN may suffice. If not, we can try a simple CNN.

Furthermore, the architecture is scalable to a large group of people. Since the GAN and classification are done on a person, it can be easily scaled up for a large group of people.

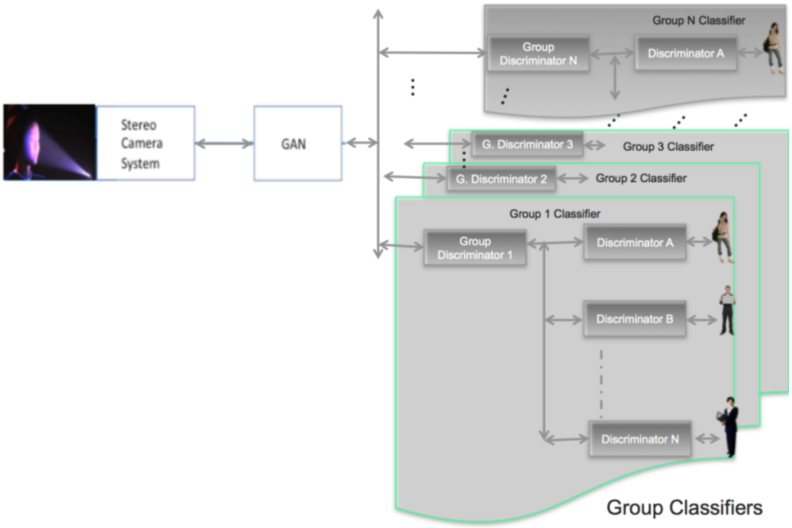


**Fig. 3** Personalized GAN Trained Group Classifiers

**WAN-GAN Architecture relative to DCGAN**

The discriminator accuracy improved a lot for example

| | |
|---|---|
| Alex Net *Reduce error from*26.2% | - 15.4 % |
| ZF Net | 11.2 % |
| VGG | 7.3% |
| Google | 6.7% |
| Res Net | 3.6% |

While human Error 10-5 % it is evident that the facial recognition will have a lot of application and it is under utilized. In our experiment we have the same discriminator used at two layers, one as part of the GAN and the second as a discriminator.

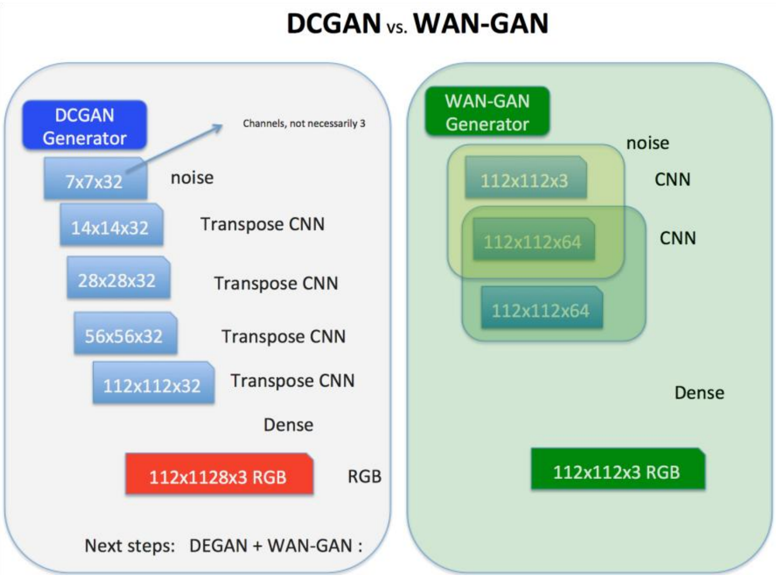


**Fig. 4** WAN-GAN Discriminator relative to VGG and other Source

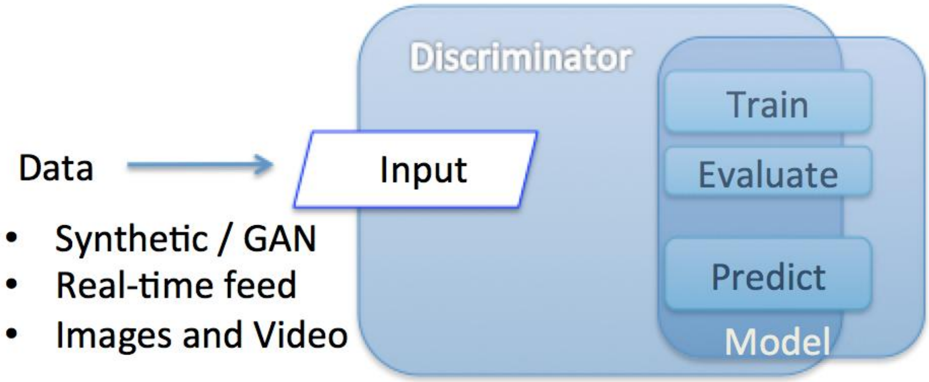


**Fig. 5**

While VGG states 91-92% accuracy on being the target or not, our research achieved 81 % on provided the accuracy on scale 1 to 10, 81% on this scale is much higher than 92% on scale true or false. The reason this is significant improvement, as facial recognition create a scale for security and guarantee accuracy based on desired application 50% is not good enough for security applications for example.

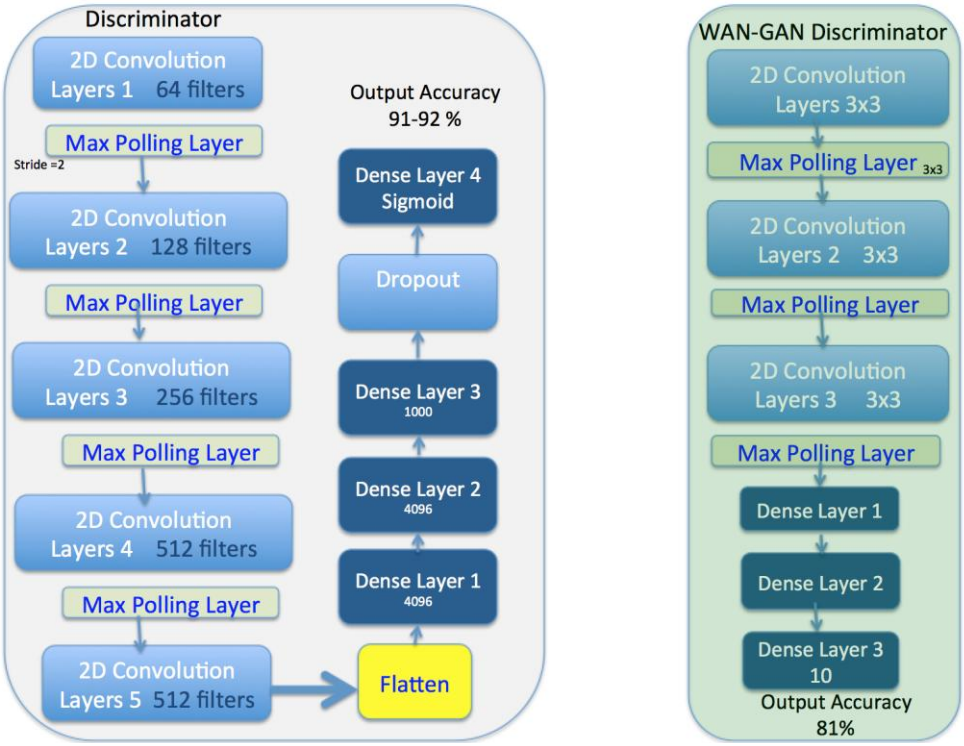


**Fig. 6**

https://towardsdatascience.com/first-contact-with-tensorflow-estimator-69a5e072998d
Source :https://torres.ai/

**Online Document**

[6] J Cartwright, Big stars have weather too, *IOP Publishing Physics*, Web. http://physicsweb.org/articles/news/11/6/16/1, **2007**.

Since most of your references are online documents, use the above format to report them.

## REFERENCES

[1]. J. Brownlee, "How to Code the GAN Training Algorithm and Loss Functions," *Machine Learning Mastery*, July 12, 2019. [Online]. Available: https://machinelearningmastery.com/how-to-code-the-generative-adversarial-network-training-algorithm-and-loss-functions///. [Accessed: Nov. 5, 2019].

[2]. O.Kharkovyna, "An Intro to Deep Learning for Face Recognition," *mc.ai*, June 26, 2019. [Online]. Available: https://mc.ai/an-intro-to-deep-learning-for-face-recognition/. [Accessed: Nov. 5, 2019].

[3]. J. Brownlee, "How to Use the Up sampling 2D and Conv2D Transpose Layers in Keras," *Machine Learning Mastery*, June 24, 2019. [Online]. Available: https://machinelearningmastery.com/upsampling-and-transpose-convolution-layers-for-generative-adversarial-networks//. [Accessed: Nov. 5, 2019].

[4]. J. Brownlee, "Tips for Training Stable Generative Adversarial Networks," *Machine Learning Mastery*, June 19, 2019. [Online]. Available: https://machinelearningmastery.com/how-to-train-stable-generative-adversarial-networks//. [Accessed: Nov. 5, 2019].

[5]. Z. Wang, et al., "Generative Adversarial Networks: A Survey and Taxonomy," *Insight Centre for Data Analytics*, Dublin City University, Dublin, Ireland, June 14, 2019. [Online]. Available: https://arxiv.org/pdf/1906.01529.pdf. [Accessed: Nov. 5, 2019].

[6]. M. Li, et al., "D2PGGAN: Two Discriminators Used in Progressive Growing of GANs," *University of Science & Technology of China*, May 2019. [Online]. Available: https://www.researchgate.net/publication/332790169_D2PGGAN_Two_Discriminators_Used_in_Progressive_Growing_of_GANS. [Accessed: Nov. 5, 2019].

[7]. M.Peixeiro, "Everything You Need to Know About Autoencoders in Tensorflow: From Theory to Implementation in TensorFlow," *Towards Data Science*, April 30, 2019. [Online]. Available: https://towardsdatascience.com/everything-you-need-to-know-about-autoencoders-in-tensorflow-b6a63e8255f0. [Accessed: Nov. 5, 2019].

[8]. J-Y.Baek, et al., "Adversarial Learning with Knowledge of Image Classification for Improving GANS,"*IEEE Access*, April 29, 2019. [Online]. Available: https://www.semanticscholar.org/paper/Adversarial-Learning-With-Knowledge-of-Image-for-Baek-Yoo/264f825da08b1f3d9cace68d7df8bc1606898876. [Accessed: Nov. 5, 2019].

[9]. A. Elmahmudi, et al. "Deep Face Recognition Using Imperfect Facial Data," *Future Generation Computer Systems*, 99 (2019) 213-225. [Online]. Available: https://reader.elsevier.com/reader/sd/pii/S0167739X18331133?token=1376129E5BD3D0BB737F2258651A736619CE7E81AD8E2BB73B69C69C80144EF251F019B4F180A6CAF82307C268E0B118. [Accessed: Nov. 5, 2019].

[10]. K. Ahirwar, et al. "The Rise of Generative Adversarial Networks," *KDNuggets*, April 2019. [Online]. Available: https://www.kdnuggets.com/2019/04/rise-generative-adversarial-networks.html. [Accessed: Nov. 5, 2019].

[11]. S. Tariq, et al. "GAN is a Friend or a Foe: a Framework to Detect Various Fake Face Images," *State University of New York*, SUNY-Korea, Incheon, South Korea, April 2019. [Online]. Available: https://www.researchgate.net/publication/332812263_GAN_is_a_friend_or_foe_a_framework_to_detect_various_fake_face_images. [Accessed: Nov. 5, 2019].

[12]. P. Lei, et al. "On Low-Resolution Face Recognition in the Wild: Comparisons and New Techniques," 3/28/19. [Online]. Available: https://arxiv.org/pdf/1805.11529.pdf. [Accessed: Nov. 5, 2019].

[13]. "Fathers of the Deep Learning Revolution Receive ACM A.M. Turing Award," *Association for Computing Machinery*, March 27, 2019. [Online]. Available: https://www.acm.org/media-center/2019/march/turing-award-2018. [Accessed: Nov. 5, 2019].

[14]. J.D. Curto, et al. "High-Resolution Deep Convolutional Generative Adversarial Networks," 3/24/19. [Online]. Available: https://arxiv.org/pdf/1711.06491.pdf. [Accessed: Nov. 5, 2019].

[15]. C. Shorten. "Deeper Into DCGANs," *Towards Data Science*, 2/13/19. [Online]. Available: https://towardsdatascience.com/deeper-into-dcgans-2556dbd0baac. [Accessed: Nov. 5, 2019].

[16]. M. Wang, et al. "Deep Face Recognition: A Survey," *School of Information and Communication Engineering*, Beijing University of Posts and Telecommunications, Beijing, China, 2/12/19. [Online]. Available: https://arxiv.org/pdf/1804.06655.pdf. [Accessed: Nov. 5, 2019].

[17]. C. Hardy, et al. "Multi-Discriminator Generative Adversarial Networks for Distributed Datasets,"*INRIA*, 2/6/19. [Online]. Available: https://arxiv.org/abs/1811.03850?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%253A+arxiv%252FQSXk+%2528ExcitingAds%2521+cs+updates+on+arXiv.org%2529. [Accessed: Nov. 5, 2019].

[18]. A. Bonner. "The Complete Beginner's Guide to Deep Learning: Convolutional Neural Networks and Image Classification," *Towards Data Science, 2/2/19.* [Online]. Available: https://towardsdatascience.com/wtf-is-image-classification-8e78a8235acb. [Accessed: Nov. 6, 2019].

[19]. J. Vincent. "These Faces Show how Far AI Image Generation has Advanced in Just Four Years," *The Verge*, 12/17/18. [Online]. Available: https://www.theverge.com/2018/12/17/18144356/ai-image-generation-fake-faces-people-nvidia-generative-adversarial-networks-gans.[Accessed: Nov. 5, 2019].

[20]. S. Wolf. "ProGAN: How NVIDIA Generated Images of Unprecedented Quality,"*Towards Data Science*, 12/16/18. [Online]. Available: https://towardsdatascience.com/progan-how-nvidia-generated-images-of-unprecedented-quality-51c98ec2cbd2.[Accessed: Nov. 5, 2019].

[21]. J-Y. Zhu, et al. "Generative Visual Manipulation on the Natural Image Manifold," in European Conference on Computer Vision (ECCV), 12/16/18 (2016). [Online]. Available:

https://arxiv.org/pdf/1609.03552.pdf. [Accessed: Nov. 5, 2019].

[22]. S Saha, "A Comprehensive Guide to Convolutional Neural Networks – The ELI5 Way," *Towards Data Science*, 12/15/18. Available at: https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53. [Accessed: Nov. 6, 2019].

[23]. S. Wolf. "CycleGAN: Learning to Translate Images (Without Paired Training Data),"*Towards Data Science*, 11/19/18. [Online]. Available: https://towardsdatascience.com/cyclegan-learning-to-translate-images-without-paired-training-data-5b4e93862c8d. [Accessed: Nov. 6, 2019].

[24]. T. Lane. "Transposed Convolutions Explained with MS Excel!" *MXNet*, 11/2/18. [Online]. Available: https://medium.com/apache-mxnet/transposed-convolutions-explained-with-ms-excel-52d13030c7e8. [Accessed: Nov. 6, 2019].

[25]. Z. Lin, et al. "PacGAN: The Power of Two Samples in Generative Adversarial Networks," 32nd Conference on Neural Information Processing Systems, 11/2/18, (NeurIPS 2018), Montreal, Canada. https://papers.nips.cc/paper/7423-pacgan-the-power-of-two-samples-in-generative-adversarial-networks.pdf

[26]. B. Marr. "What is Deep Learning in AI? A Simple Guide with 8 Practical Examples," *Forbes*, October 1, 2018. [Online]. Available: https://www.forbes.com/sites/bernardmarr/2018/10/01/what-is-deep-learning-ai-a-simple-guide-with-8-practical-examples/#6de589d18d4b. [Accessed: Nov. 6, 2019].

[27]. J. Zeng, et al. "Deep Convolutional Neural Network Used in Single Sample per Person Face Recognition," *Computational Intelligence and Neuroscience*, Volume 2018, Article ID 3803627, 11 pages, 8/23/18. [Online]. Available: https://www.researchgate.net/publication/327194037_Deep_Convolutional_Neural_Network_Used_in_Single_Sample_per_Person_Face_Recognition. [Accessed: Nov. 6, 2019].

[28]. M. Aspris. "Is the Transposed Convolution Layer and Convolution Layer the Same Thing? Experimenting with Concepts Using PyTorch,"*Towards Data Science*, 7/17/18. [Online]. Available: https://towardsdatascience.com/is-the-transposed-convolution-layer-and-convolution-layer-the-same-thing-8655b751c3a1. [Accessed: Nov. 6, 2019].

[29]. J. Hui. "GAN – DCGAN (Deep Convolutional Generative Adversarial Networks)," *Medium*, 6/17/18. [Online]. Available: https://medium.com/@jonathan_hui/gan-dcgan-deep-convolutional-generative-adversarial-networks-df855c438f. [Accessed: Nov. 6, 2019].

[30]. J. Hui. "GAN – Super Resolution GAN (SRGAN)," *Medium*, 6/14/18. [Online]. Available: https://medium.com/@jonathan_hui/gan-super-resolution-gan-srgan-b471da7270ec. [Accessed: Nov. 6, 2019].

[31]. J. Hui. "GAN – Unrolled GAN (how to reduce mode collapse)," *Medium*, 6/9/18. [Online]. Available: https://medium.com/@jonathan_hui/gan-unrolled-gan-how-to-reduce-mode-collapse-af5f2f7b51cd. [Accessed: Nov. 6, 2019].

[32]. K. Aggarwal. "Transpose Convolution," *Towards Data Science*, 5/19/18. [Online]. Available: https://towardsdatascience.com/transpose-convolution-77818e55a123. [Accessed: Nov. 6, 2019].

[33]. D. Mwiti. "Convolutional Neural Networks: An Intro Tutorial," Heartbeat / Medium, 5/8/18. https://heartbeat.fritz.ai/a-beginners-guide-to-convolutional-neural-networks-cnn-cf26c5ee17ed. [Accessed: Nov. 6, 2019].

[34]. K. Eykholt, et al. "Robust Physical-World Attacks on Deep Learning Visual Classification,"*Computer Vision and Pattern Recognition (CVPR) 2018*,4/10/18. [Online]. Available: https://arxiv.org/pdf/1707.08945.pdf. [Accessed: Nov. 6, 2019].

[35]. Y. Wang. "A Fully Progressive Approach to Single-Image Super Resolution," *2018 IEEE/Computer Vision and Pattern Recognition (CVPR)*, 4/10/18. [Online]. Available: https://arxiv.org/abs/1804.02900. [Accessed: Nov. 6, 2019].

[36]. S. Arora, et al. "Limitations of Encoder-Decoder GAN Architectures," *Off the Convex Path*, 3/12/18. [Online]. Available: https://www.offconvex.org/2018/03/12/bigan/. [Accessed: Nov. 6, 2019].

[37]. T. Karras, et al. "Progressive Growing of GANs for Improved Quality, Stability, and Variation," *International Conference on Learning Representation (ICLR) 2018*, 2/26/18.[Online]. Available: https://arxiv.org/abs/1710.10196. [Accessed: Nov. 6, 2019]

[38]. J-D. Seo. "Only Numpy: Understanding Back Propagation for Transpose Convolution in Multi-Layer CNN with Example and Interactive Code," *Towards Data Science*, 1/29/18. [Online]. Available: https://towardsdatascience.com/only-numpy-understanding-back-propagation-for-transpose-convolution-in-multi-layer-cnn-with-c0a07d191981. [Accessed: Nov. 6, 2019].

[39]. "Ian Goodfellow: Generative Adversarial Networks (NIPS 2016 Tutorial)," 1/18/18. https://www.youtube.com/watch?v=HGYYEUSm-0Q

[40]. "A Guide to Convolution Arithmetic for Deep Learning," by Vincent Dumoulin and Francesco Visin, 1/12/18. https://arxiv.org/pdf/1603.07285.pdf

[41]. "On the Tradeoff Between Mode Collapse and Sample Quality in Generative Adversarial Networks," 2018 IEEE Conference on Signal and Information Processing (Global SIP).

https://www.researchgate.net/publication/331850923_ON_THE_TRADEOFF_BETWEEN_MODE_COLLAPSE_AND_SAMPLE_QUALITY_IN_GENERATIVE_ADVERSARIAL_NETWORKS

[42]. "To Learn Image Super Resolution, Use a GAN to Learn How to Do Image Degradation First," by Adrian Bulat, et al., Computer Vision Laboratory, University of Nottingham, United Kingdom, ECCV 2018. https://www.adrianbulat.com/downloads/ECCV18/image-super-resolution.pdf

[43]. "High Resolution Image Synthesis and Semantic Manipulation with Conditional GANs," by Ting Chun Wang, et al., CVF 2018. http://openaccess.thecvf.com/content_cvpr_2018/papers/Wang_High-Resolution_Image_Synthesis_CVPR_2018_paper.pdf https://github.com/NVIDIA/pix2pixHD https://github.com/aleju/papers/blob/master/neural-nets/High_Resolution_Image_Synthesis_with_Conditional_GANs.md

[44]. "Up-Sampling with Transposed Convolution," by Naoki Shibuya, Medium Corporation, 11/13/17. https://medium.com/activating-robotic-minds/up-sampling-with-transposed-convolution-9ae4f2df52d0

[45]. "Strengths and Weaknesses of Deep Learning Models for Face Recognition against Image Degradation," KlemenGrm, et al., Submitted for publication to IET Biometrics, 10/4/17. https://arxiv.org/pdf/1710.01494.pdf

[46]. "Dual Discriminator Generative Adversarial Nets," by TuDinh Nguyen, Trung Le, et al., 31st Conference on Neural Information Processing Systems, (NIPS 9/12/17). https://arxiv.org/pdf/1709.03831.pdf

[47]. "Generative Adversarial Networks for Improving Face Classification," by Jonas Natten, University of Agder, September 2017. https://uia.brage.unit.no/uia-xmlui/bitstream/handle/11250/2454822/Natten%2c%20Jonas.pdf?sequence=1&isAllowed=y

[48]. "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network," by Christian Ledig, et al., CVPR 5/25/17. https://arxiv.org/abs/1609.04802

[49]. "Unrolled Generative Adversarial Networks," ICLR 5/12/17 (last revised). https://arxiv.org/abs/1611.02163

[50]. "Began: State of the Art Generation of Faces with Generative Adversarial Networks," Heuritech Tech Blog, 4/11/17. https://lab.heuritech.com/began-state-of-the-art-generation-of-faces-with-generative-adversarial-networks

[51]. "GP-GAN: Towards Realistic High-Resolution Image Blending," Semantic Scholar, 3/25/17. https://arxiv.org/abs/1703.07195, https://github.com/wuhuikai/GP-GAN

[52]. "Unrolled Generative Adversarial Networks," by Luke Metz, Google Brain, et al., 1/15/17. https://arxiv.org/pdf/1611.02163v2.pdf

[53]. "Super Resolution with Generative Adversarial Networks," by Boris Kovalenko, Stanford, California, 2017. http://cs231n.stanford.edu/reports/2017/pdfs/17.pdf

[54]. "Facial Expression Recognition Using Convolutional Neural Networks: State if the Art," by Christopher Pramerdorfer, et al., Computer Vision Lab, TU Wien, Vienna, Austria, 12/9/16. https://arxiv.org/pdf/1612.02903.pdf

[55]. Ian Goodfellows' PhD Defense, 9/3/14. https://www.youtube.com/watch?v=ckoD_bE8Bhs

[56]. "Generative Adversarial Networks," Ian J. Goodfellow, et al., Cornell University, June 10, 2014. https://arxiv.org/abs/1406.2661