Vit-GAN: Image-to-image Translation with Vision Transformes and Conditional GANS

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

11-10-2021 / 20-10-2021

CITATION

Gunduc, Yigit (2021): Vit-GAN: Image-to-image Translation with Vision Transformes and Conditional GANS. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.16785751.v1>

DOI

[10.36227/techrxiv.16785751.v1](https://doi.org/10.36227/techrxiv.16785751.v1)

Vit-GAN: Image-to-image Translation with Vision Transformers and Conditional GANS

Yiğit Gündüç¹

¹ygunduc@gmail.com

October 11, 2021

Abstract

In this paper, we have developed a general-purpose architecture, Vit-Gan, capable of performing most of the image-to-image translation tasks from semantic image segmentation to single image depth perception. This paper is a follow-up paper, an extension of generator based model [1] in which the obtained results were very promising. This opened the possibility of further improvements with adversarial architecture. We used a unique vision transformers-based generator architecture and Conditional GANs(cGANs) with a Markovian Discriminator (PatchGAN) (<https://github.com/YigitGunduc/vit-gan>). In the present work, we use images as conditioning arguments. It is observed that the obtained results are more realistic than the commonly used architectures.

1 Introduction

Most of the problems we face in computer vision can directly or indirectly benefit from image to image translation. The image-to-image translation problem is mapping an input image to a corresponding output image. Image-to-Image translation has a wide range of applications, such as object transfiguration, season transfer, image segmentation, photo enhancement, or in general, where an image is needed to be mapped onto another. Most of the traditional translation methods are only applicable to a single domain or trained for a special task [6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. In this work, we propose a general-purpose architecture that is capable of performing most image-to-image translation tasks, from semantic image segmentation to single image depth precision.

Optimization and loss metric is one of the biggest problems in deep learning. The same problem also exists in the Image-to-Image translation. In our experiment, we see that the models, which are optimized with the Mean Absolute Error or Mean Square Error functions tend to generate blurry results. This problem is related with the optimization procedure which averages over the values. Hence such a model is unable to perform well on complex tasks. GANs [17] are the perfect fit for sharp and realistic images as output. In a GAN model, the main idea is; the discriminator learns to classify if the output image is real or fake, and the generator tries to fool it. During simultaneous training, both components play a min-max game and minimize their losses. Generators in GAN models do not learn how to generate an image but

rather the loss function itself the discriminator. GANs classify images as real or fake. That makes them powerful tools for the generation of images in as close resemblance as the real ones.

The proposed solution applies to nearly all image-to-image translation tasks and problems with the help of a relatively new concept of transformers, specifically vision transformers. In this work, we have explored how vision transformers can be used for image-to-image translation. In the current work, we used CNNs alongside patches, patch encoder, and transformers. In this paper, we introduced a unique vision transformers-based generator and Conditional GANs(cGANs) [2] with a Markovian discriminator (PatchGAN) [3]. This discriminator has been applied to similar problems and shown promising results [4]. We also used the PatchGAN discriminator, which is applied to reinforce our generator to generate sharp and accurate images on difficult cases.

2 Background

In the literature, a variety of algorithms proposed for image-to-image translation [5] (and references therein). The majority of the methods were based on CNNs and employed encoder-decoder networks [6] for image-to-image translation. In an autoencoder architecture, the input image passes through the encoder. The encoder layers downsample the input until a bottleneck layer. In a reversed process, starting from the bottleneck point, the decoder upsamples the image until the original shape is restored. The problem with encoder-decoder architectures is that low-level information is lost while the input images pass through the bottleneck. The information loss may have negative effects while constructing the output image. Some successful methods also utilized U-Net [18] like approaches to overcome this problem. U-Net adds skip connections between layers. if one wants to formulate skip connections between the i^{th} and the $(n - i)^{\text{th}}$ layers, where n is the layer number, one need to create a link in between which acts as attention.

For image-to-image translation, Generative Adversarial Network [19] (GAN) based approaches also have been used. A GAN is a generative model, which generates new instances of data used for training. A GAN [19] consist of two-part a generator and a discriminator. The discriminator is used to lead the generator the generate shapes and accurate images. These networks play a min-max game against each other to minimize their losses. Through the competition, GANs can generate realistic output images from random noise input. Conditional generative adversarial network [2] is a GAN architecture that enables the generation of output images according to introduced conditions. The condition can be a class label or any other descriptive information on the dataset.

Vaswani et al. proposed Transformer architectures for machine translation [20]. Since then, the Transformers became the state of the art in most of the NLP tasks. The Transformer models are often pre-trained on large text corpora and then fine-tuned for specific tasks [21, 22, 23]. Transformers, with their unique self-attention mechanisms and MLPs, do not suffer from memory shortages. Comparing with the Long Short-Term Memory Networks (LSTMs), which cannot recall tokens after a certain amount of time, transformers show their advantages. There have been attempts to utilize transformers like self-attention architectures with CNNs for computer vision tasks [24, 25, 26, 27, 28].

Image GPT(iGPT) [29] is another notable architecture that takes pixel sequences as inputs and tries to

generate output. Unlike vision Transformers, where they take all the images and split them into patches, iGPT treats pixels just like tokens of words. iGPT takes some pixels and outputs some pixels similar to how a text generation model would work.

Vision transformers (Vit) [9] is a transformer architecture that only uses transformer encoders and a custom patch embedding to work with images. Vision Transformers are used for image classification and reached an accuracy of 88.55% on ImageNet. Vit works by splitting an image into patches, flattens those patches, produces lower-dimensional linear embeddings from the flattened patches, adding the positional embeddings, and then feeding those sequences as an input to a transformer encoder. Vit models are usually pre-trained on a large dataset and then finetuned on smaller datasets for image classification.

3 Method

3.1 Overview

The most significant difference between our approach and the other image-to-image translation methods is its unique generator: A hybrid architecture that utilizes vision transformers alongside convolutional layers to convert images [1]. Unlike prior methods, the proposed method utilizes vision transformers for image understanding and uses PatchGAN [3] to ensure generated results are clear, accurate, and sharp. In the present work, we use images as conditioning arguments. The generator uses the condition input image to make a selective prediction. The discriminator takes the output of the generator as input. The discriminator output and L1 loss are used in conjunction to compute the loss. Lambda, the coefficient which accompanied the L1 loss, is chosen as 100. This coefficient can be tuned if the model generates visual artifacts. This situation never happened in our experiments. For complex tasks, a cGAN alongside L1 loss gives better image quality [4]. Finally, the optimizer uses the total loss for applying gradients to the weights. Please see, <https://github.com/YigitGunduc/vit-gan> for details of the model.

3.2 Patches and Patch Encoding

Our generator uses Vit-like patches and patch embedding to understand the image. The model takes inputs, splits them into patches. The number of patches depends on the patch size and image size (image height $\times$ image width). The formula $(Imageheight/patchsize) * *2$ (assuming image height and width are the same) is used to calculate the number of patches. Then, all patches are flattened and reshaped into $1 - D$ sequences of patches. We use a trainable embedding layer to embed patches to latent space vectors and add a learnable $1 - D$ position embedding to the patches. This last embedding preserves the positions of the patches in the image.

3.3 Transformer Layers

The Generator consists of N identical stacked layers [1]. Each stacked-layer consists of a multi-head self-attention (MHA) trailed by a feed-forward linear layer. MHA contains H heads where each one of the identical heads computes independent scaled dot-product attention. Therefore the model can attend different places of the embedding vector. Then transformation is applied to concatenate the attention results obtained from heads before further processing in the model.

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O$$

The utilized attention is the same as used in the original paper [20].

The proposed and formula is as follows:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where, Q , K , V are the query, key, and the value respectively.

After the dot-product attention is applied we scale the outputs since the dot-product attention tends to output large values. After the normalization process, finalize the outputs with a feed-forward linear layer.

3.4 Architecture

3.4.1 Generator

The generator section of the architecture is taken from a previous paper on image-to-image translation. The success of the generator on image translation has already been proven. [1]. Inputs of the generator are introduced to the patch encoders. After images are split into patches and passed through the patch encoder, the resulting vectors are input to the transformer encoders. N of the transformer encoders are stacked on top of each other, which are used to have an understanding of the context. In this image-to-image translation task, this context is the image. The transformer encoder's output is passed to a residual block. The residual block consists of two convolutional layers with skip connections in between. ReLU and batch normalization are applied in the residual blocks. Residual blocks are followed by an upsampling block. Transpose convolutions are trailed by LeakyReLU and batch normalization in the upsampling block.

The experiments with various generator architectures are carried out and it is concluded that using a combination of Residual blocks and transpose convolutions gives the best output [1].

3.4.2 Discriminator (PatchGAN)

A carefully chosen loss function plays the utmost importance in the performance of the model. pix2pix [4] paper has experimented on this topic and shown that L1 loss with a cGAN outperforms better than the other methods. Hence, we applied the same optimization strategy in our model with L1 loss alongside a cGAN.

If the discriminator is not conditional, the model tends to generate results that fool the discriminator but, the output images resemble any one of the training data set. In other words, the model treats the input image as a random noise like most of the GANs algorithms do. This situation is not acceptable for image-to-image translation tasks. The output must be built upon the input and share the same low-level features.

PatchGAN is a discriminator for GANs which penalizes local image patches rather than the whole image. Discriminator tries to classify $N \times N$ patches as real or fake. Classifying patches rather than the whole image enables PatchGAN to run and train much faster. In our experimentation with the model, we also concluded that PatchGAN outperforms traditional discriminators. Hence PatchGAN is selected as the discriminator for our model.

4 Experiments

We tested our method on a different task to see how well it performed. We used image segmentation, depth prediction from a single image, architectural labels to photos. After training, the model performed excellently. We concluded that such a performance is related to the data set.

4.1 Datasets

We have tested our method on three different datasets for three different Tasks. Cityscapes [30, 31], Facades [32], and RGB-D [33] datasets are used to test, models prediction ability of semantic segmentation, convert architectural labels to facades and depth mapping from a single image.

Cityscapes [30, 31] is a large-scale database that is explicitly built for semantic segmentation. The database contains urban street scenes which provide semantic instances of 30 classes grouped into 8 categories (flat surfaces, humans, vehicles, constructions, objects, nature, sky, and void). The dataset consists of around 5000 annotated images. Images from 50 cities in Germany with different lighting and weather conditions constitute the dataset. We trained our model by using this dataset for both ways, to generate semantic segmentation maps (Figure 4) when the photo is inputted and generate a photo of the scene when a segmentation map has been inputted (Figure 5).

Facades [32] is a database of facade images with its corresponding labels of machine perception. The dataset includes 606 annotated rectified images from diverse parts of the world with distinct architectural

Figure 1: Facades dataset results generated by custom discriminator (used transformers).

styles of facades obtained from different sources.

Our model has been trained on the Facades dataset to convert architectural labels to facades (Figure 6).

RGB-D [33] dataset consists of various indoor and outdoor images with associated depth maps. Indoor and outdoor samples are taken from offices, rooms, dormitories, exhibition centers, streets, the road from Yonsei University and Ewha University. The model is trained on the RGB-D dataset to experiment with the depth prediction precision from a single image. The model successfully generates depth maps from a single raw input image (Figure 7).

4.2 Experiment on Custom Discriminator

We also experimented with a custom PatchGAN like discriminator architecture of our creation. The proposed algorithm utilizes transformer layers instead of convolutional layers. The algorithm works by splitting the image into patches, generating patches linear embedding, adding position embeddings to preserve the positional information. Transformer Encoders use the embedded patches as input. Finally, we used a Convolutional layer to classify $N \times N$ patches as real or fake, here N is the width and height of the last convolutional layer outputs. In this architecture, the convolutional layer stands for matching input and output shapes. Instead of a convolutional layer, up or downsampling of the transformer encoder’s output is also possible. This approach does not work as well as the convolutional layer but is still usable and relatively well. The biggest drawback of not using a convolutional layer is that the generated image is pixelated and not realistic as we want them to be Figure 1 shows the result of using the experimental discriminator in the architecture.

The traditional PatchGAN discriminator, made out of convolution layers, was used to obtain displayed results throughout the present paper.

4.3 Image Quality wise Comparison with Other Architectures

Evaluation of the performance of different models is not an easy task. A simple Mean Squared Error or cross-entropy is not sufficient to decide which models outperform the others. In this work, we used some of the most advanced image evaluation metrics FID(Frechet Inception Distance) [19], SSIM(structural similarity) [34], IS(inceptions score) [35]. Inceptions score is one of the widely used evaluation metrics for GANs and is in most experiments. We compared our architecture with U-Net and autoencoders that are the go-to choices for most image-to-image translation tasks such as image segmentation. We trained an autoencoder, a U-Net, and our approach on the cityscapes dataset to convert segmentation maps to photos. Table 1 shows the measures of the generated image quality among the different architectures according to the metrics we mentioned above. For the visual comparison, figure 2 presents outputs of all three models.

Model	FID	IS	SSIM
Tensor-to-image(Ours)	939	1.281	0.46
U-Net	4318	1.280	0.36
Autoencoder	6434	1.343	0.25

Table 1: Performance comparissons

Results, shown in figure 2 and table 1, indicate that the autoencoder fails to perform as well as the competitors. It has a poor and low-level understanding of the image. Autoencoder architecture is only able to distinguish the ground from the sky, thus generates unrealistic images. The generated images mostly contain two colors, one for the sky and the other for the land. U-Net is an autoencoder with skip connections between encoder and decoder layers. U-Net was able to pick up the different objects in the scene and fill them with different colors. Despite the performance increase, the generated images were not realistic enough. They contained the contours of objects with different colors without shadows or reflections. Correct identification colors, shadows, and reflections are the main elements that make an image closer to the photo-realistic level.

4.4 Contribution of Adversary

In a previous paper [1] we have used a non-adversarial model which was trained with only a custom-designed generator. The results were fascinating and promising, which opened the possibility of further improvements with adversarial architecture. We have decided to extend its domain by adding a discriminator to the network and generator. We have trained both models on the facades dataset to validate the generation quality of the output. The expected is the generation of photos of the building just by the architectural labels. The models exhibit slight differences in training. The adversarial model is trained with a cGAN alongside the L1 loss. The non-adversarial model (only generator) is optimized against the L1 loss. The model trained with a cGAN was quite successful in generating sharp and realistic images. The basic structures of the building, windows, doors, and edges, were in the right places. The model optimized against the L1 loss has only a basic understanding of the shape of buildings. Also, the generated images

are blurry and not clear. So we have concluded that detailed image-to-image translation tasks require the adversary model for better performance (Figure 3).

4.5 Object segmentation

Object segmentation is more of a vision problem than it is an image-to-image translation. It is all about finding the appropriate regions and filling them with corresponding colors. But since both the input and the output are images, we used our model to segment objects. We trained on the Cityscapes [30, 31] dataset. The raw images are used as input and expected the model to output segmentation maps. It is a difficult task for a model to accomplish without a discriminator. Hence we presented all the object segmentation results on this paper using the model trained with a discriminator (Figure 4).

4.6 Image-to-image translation

Although we trained our model on different tasks, we all did them in image-to-image translation fashion where we input a 3D tensor and expected a 3D or 1D tensor. A generator-only approach is suitable for most of the tasks. Nevertheless, the model which contains a discriminator is a better solution for the close similarity between the output image and the dataset. Discriminator forces the generator to generate results that are sharp and realistic. Hence using a discriminator is the go-to choice (Figure 3).

4.7 Depth estimation from a single image

Estimating the depth of an image is an important problem and has a wide range of applications like self-driving cars, drones, and more. Often depth prediction is expensive because it requires specialized hardware and multiple shots of every scene. With the proposed model, we aim to generate depth maps from only one single image. This approach is economical and successful application is superior comparing with the classical two shut system. It reduces the cost of accomplishing the depth map of an image and increases the accessibility of the depth maps. Shuts of only one camera are enough to estimate depth in our approach (Figure 7).

5 Conclusion

The traditional image-to-image translation methods are only applicable to a single domain or trained for predefined tasks. In this work, We used a unique vision transformers-based generator architecture to extend the range of vision tasks with a single architecture. This paper is an extension of previous work, generator-based model, Tensor-to-Image [1]. The proposed architecture uses Conditional GANs(cGANs) with a Markovian discriminator (PatchGAN). Adding a discriminator into the architecture [1] improved the output image quality and the range of vision tasks. Comparing with the commonly used models,

namely, autoencoder and U-Net architectures, the proposed method shows better performance on the image-to-image translation tasks.

We concluded that vision transformers are capable of performing most of the image-to-image translation tasks. Better image quality is reinforced with an architecture containing a discriminator. The model details can be found in GitHub account (<https://github.com/YigitGunduc/vit-gan>).

Figure 2: Comparison of three different architectures. Photos are, left to right, input, target, autoencoder output, U-Net output, and output obtained using our method.

Figure 3: Visual Comparison cGan and No cGan tests

Figure 4: Generation of semantic segmentation map

Figure 5: Generation of a photo of the scene when a segmentation map has been inputted

Figure 6: Facets dataset: Converting architectural labels to facades

Figure 7: RGB-D Dataset: Experiment on Depth Perception of the Model

References

- [1] Y. Gündüç, Tensor-to-Image: Image-to-Image Translation with Vision Transformers, TechRxiv. Preprint, <https://doi.org/10.36227/techrxiv.16727140.v1>, (Details: <https://github.com/YigitGunduc/tensor-to-image>), 2021.
- [2] M. Mirza and S. Osindero, Conditional Generative Adversarial Nets, arXiv1411.1784[cs.LG] , 2014.
- [3] C. Li and M. Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. ECCV, 2, 4, 2016.
- [4] P. Isola, and J-Y. Zhu, and T. Zhou, and A. A. Efros, Image-to-Image Translation with Conditional Adversarial Networks, CVPR, 2017.
- [5] P. Yingxue, L. Jianxin, Q. Tao, and C. Zhibo, Image-to-Image Translation: Methods and Applications <https://arxiv.org/pdf/2101.08629.pdf> 2021.
- [6] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. In SIGGRAPH, 1, 4, 2001.
- [7] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin. Image analogies. In SIGGRAPH, 1, 4, 2001.
- [8] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. ACM Transactions on Graphics (TOG), 25(3):787– 794, 1, 2006.
- [9] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In CVPR, 1, 2005.
- [10] T. Chen, M.-M. Cheng, P. Tan, A. Shamir, and S.-M. Hu. Sketch2photo: internet image montage. ACM Transactions on Graphics (TOG), 28(5):124, 1, 2009.
- [11] Y. Shih, S. Paris, F. Durand, and W. T. Freeman. Data-driven hallucination of different times of day from a single outdoor photo. ACM Transactions on Graphics (TOG), 32(6):200, 1, 2013.
- [12] P.-Y. Laffont, Z. Ren, X. Tao, C. Qian, and J. Hays. Transient attributes for high-level understanding and editing of outdoor scenes. ACM Transactions on Graphics (TOG), 33(4):149, 1, 4, 16, 2014.
- [13] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In CVPR, 1, 2, 5, 2015.
- [14] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In ICCV, 1, 2015.
- [15] S. Xie and Z. Tu. Holistically-nested edge detection. In ICCV,1,2,4,2015.
- [16] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. ECCV, 1, 2, 5, 7, 8, 16, 2016.
- [17] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In NIPS, 2, 4, 6, 7, 2014.

- [18] O. Ronneberger , P. Fischer, T. Brox, U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab N., Hornegger J., Wells W., Frangi A. (eds) Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015. MICCAI 2015. Lecture Notes in Computer Science, vol 9351. Springer, Cham. https://doi.org/10.1007/978-3-319-24574-4_28, 2015.
- [19] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In Advances in neural information processing systems, 30, pp. 6626–6637, (arXiv:1706.08500) 2017.
- [20] A. Vaswani and N. Shazeer and N. Parmar and J. Uszkoreit and L. Jones and A. N. Gomez and L. Kaiser and I. Polosukhin, Attention Is All You Need, arXiv1706.03762[cs.CL] ,2017.
- [21] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In NAACL, 2019.
- [22] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. arXiv, 2020.
- [23] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. Improving language understanding with unsupervised learning. Technical Report, 2018.
- [24] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei. Relation networks for object detection. In CVPR, 2018.
- [25] F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf, Object-centric learning with slot attention. arXiv:2006.15055, 2020.
- [26] Y.-C. Chen, L. Li, L. Yu, A. Kholy, F. Ahmed, Z. Gan, Y. Cheng, and J. Liu, UNITER: UNiversal Image-TExt Representation Learning. In ECCV, 2020.
- [27] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. VisualBERT: A Simple and Performant Baseline for Vision and Language. In Arxiv, 2019.
- [28] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In ICCV, 2019.
- [29] M. Chen, and A. Radford, and R. Child, and J. Wu, and H. Jun, and P. Dhariwal, and D. Luan, and I. Sutskever, Generative Pretraining from Pixels, 2020.
- [30] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The Cityscapes Dataset for Semantic Urban Scene Understanding,” in Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [31] M. Cordts, M. Omran, S. Ramos, T. Scharwächter, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The Cityscapes Dataset,” in CVPR Workshop on The Future of Datasets in Vision, 2015.
- [32] T. Raddim Tyleček and Š. Radim, Spatial Pattern Templates for Recognition of Objects with Regular Structure, Proc. GCPR,

- [33] J. Sturm and N. Engelhard and F. Endres and W. Burgard and D. Cremers, A Benchmark for the Evaluation of RGB-D SLAM Systems, Proc. of the International Conference on Intelligent Robot Systems (IROS), 2012.
- [34] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: From error visibility to structural similarity,” *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [35] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pp. 2234–2242, 2016.