Investigating the performance of Generative Adversarial Networks on Fabric Pattern Generation

Luuk Haarman[1]
[1] *TU Delft*

Abstract

Generative adversarial networks (GANs) are a popular method for image generation, capable of state-of-the-art. Despite this popularity, the generation of fabric patterns remains somewhat unexplored. A possible reason for this could be that there is no publicly shared dataset large enough to train models. However, research into this topic could be relevant because designers can take up to multiple weeks for a single design, but a trained GAN could generate thousands of designs per day. Next to this fabric patterns have characteristics that make them different from the datasets commonly used right now, which means an investigation could yield new insights into the capabilities of GANs. This research investigated the performance of several GAN models on the task of fabric pattern generation and compared these to another popular method of image generation: variational autoencoders. By using open-source implementations of GAN models, and a dataset of fabric patterns provided by Vlisco, a Dutch fabric company, the performance of GANs has been evaluated. To protect the interests of Vlisco, this dataset cannot be shared. The results show clear promise in the capabilities of GANs to replicate these complex patterns, but these results are not yet of the same quality that GANs have been able to achieve in other domains. However, these results are an improvement from those of previous research. The results achieved by variational autoencoders are similar to those of GANs, but state-of-the-art GANs perform better. Due to limited time and resources, only low-resolution samples have been generated for evaluation. In the future, higher-resolution images could be generated which could yield additional insights into the current capabilities of GANs in this domain.

Acknowledgements

I would like to sincerely thank my supervisor Burak Yildiz, and my responsible professor Jan van Gemert for their support and guidance during this project. I would also like to thank my fellow research project group member Just Wallage for their contributions and feedback.

1 Introduction

Generative adversarial networks[1] (GANs) are a popular method for image synthesis. Since GANs were introduced many variants of GANs have been discussed and researched[4, 6, 10, 12, 14], with each model improving the performance of GANs in some way. Another research topic concerns training GANs with a relatively small dataset as GANs usually require larger datasets to get decent results, but a larger dataset is not always available. While GANs have been able to achieve positive results in many domains, other domains remain largely unexplored.

This paper will outline and discuss the results of applying several variants of GANs to a real-life problem: fabric pattern generation. While this domain is currently largely unexplored, this does not mean that it is irrelevant. These results could provide insight into the current capabilities of GANs on a domain that is different from other popular datasets such as MNIST and CIFAR-10. A possible reason for the lack of research into this domain could be that there is currently no publicly shared dataset of fabric patterns.

For this research, several variants of GANs will be trained on a dataset of fabric patterns provided by Vlisco, a Dutch company that produces and distributes fashion fabrics for the West and Central African markets. Fabric patterns can have a lot of disparity, and only share a few common traits with other patterns. Next to this patterns can often be symmetric or hold other properties which could be hard for GANs to replicate. The generated images will be compared between variants, as well the results of another popular method of image synthesis: variational autoencoders[2, 20].

The rest of the paper is organized as follows. First, the contributions of this research will be outlined, followed by a section discussing the research methodology, background information, and related research. The next section will discuss the evaluation method used, and share the parameters used for the GAN models. After this, the results will be shown for each model used, including variational autoencoders. In the next section, these results will be discussed and compared to each other next to the results of variational autoencoders.

2 Contribution of this research

This research will investigate the performance of GANs on the problem of fabric pattern generation. While GANs have been used for many areas with varying degrees of success, there is not much research available yet that focuses on this particular area. A reason for this is that there is no publicly available dataset large enough to train models with and get sufficient results. However, applying GANs to this domain could still be relevant. Designers can take days or weeks to come up with a single design. If GANs are capable of creating designs of similar quality, thousands of designs could be generated per day. Next to this, fabric patterns can be extremely complex and varied and have implicit properties such as symmetry that could make it hard for an unsupervised model to create convincing samples. Fabric patterns can vary significantly and only share some common characteristics with other patterns, which could make it hard for the model to learn these characteristics and create accurate samples. This investigation should provide insight into the current capabilities of GANs in this domain.

In comparison to previous research[16], this research will use multiple variants of GANs, including older models and state-of-the-art models such as StyleGAN2[14], which has been able to achieve high-quality results, next to being able to produce better results with smaller datasets using differential augmentation[15]. Since the original model used was DCGAN, this model will also be used to verify that the results achieved in this research are similar to previous research.

Using a multitude of GAN models selected from different years with each their distinctive features the results should show a clear development of the performance of GANs on the task of fabric pattern generation, including the current state-of-the-art performance. Also, these distinctive features could show that a specific feature performs well or poorly in this domain. This knowledge could be useful for the creation of domain-specific models, where a model is fitted to a specific domain to optimize its performance.

The results will also be compared to those of Variational Autoencoders[2, 20], which should show the differences in performances of these two methods, and provide additional insight into the current performance of GANs compared to other deep learning image generation methods in this domain.

3 Methodology

This section will describe the research methodology and provide background information with regards to the workings of generative adversarial networks (GANs) and showcase examples of new variations of GANs that have emerged over the years.

3.1 Research methodology

To properly evaluate the performance of GANs on the task of fabric pattern generation first there must be looked at previous research of the performance of GANs on the same or similar tasks. By looking at the development of the GANs used for image synthesis since the original model[1] an attempt will be made to choose distinctive GANs and compare the performance between each model. This way several key changes to the original GAN model can be evaluated and compared, next to showing the evolution of the capabilities of GANs for the domain of fabric pattern generation. The metrics used should make it easy to compare the results of a GAN model to those of other GANs and variational autoencoders. These certain characteristics in specific GANs can be compared to see which characteristics are optimal and show the development of the performance of GANs in this domain. Finding these optimal characteristics could allow the creation of a model optimally suited to this domain or motivate further research into these kinds of models.

3.2 Background information and related research

Generative adversarial networks[1] are generative models that are composed of two groups of one or more neural networks: a generator and a discriminator. The generator generates samples and the discriminator takes an image and classifies it as real or fake. The discriminator is trained to maximize the probability of correctly classifying a sample. The generator is trained to minimize the probability of the discriminator classifying a generated sample as fake. In the end, the generator should be capable of producing images indistinguishable from the original dataset.

Goodfellow et al. (2014) describes this interaction as a two-player minimax game with the following value function $V(G, D)$:

$$\underset{G}{min}\,\underset{D}{max} V(D, G) = \mathbb{E}_{x \sim p_{data}}(x)[log D(x)] + \mathbb{E}_{z \sim p_z}(z)[log(1 - D(G(z)))] \tag{1}$$

Where D is the discriminator, G is the generator, $D(x)$ represents the probability that x is real, and $G(z)$ creates a new sample given input noise variable z.

Early GANs had issues such unstable training[3], mode collapse[1], and vanishing gradients[9]. Training can be unstable because of GANs going back and forth between several local optima, failing to reach the globally optimal solution. Mode collapse limits the output of the generator if a discriminator gets stuck in a local equilibrium, and each epoch of the generator optimizes for the same discriminator, decreasing the output variety. Vanishing gradients occur when the discriminator severely outperforms the generator, which gives the generator too little information to improve its performance. However, improvements have been made to the performance of GANs and circumvent these obstacles.

In 2015 Radford et al.[4] introduced using deep convolution networks in GANs and these have since seen wide usage and resulted in sharper images. Arjovsky et al.[6] introduced Wasserstein-GANs (WGANs) in 2017 which uses Wasserstein distance as a loss function instead of the usual Jensen-Shannon divergence. WGANs attempt to remedy both mode collapse and vanishing gradients. In 2018 Jolicoeur-Martineau[12] argues that the discriminator's probability of classifying real data as real should decrease as the probability of classifying fake data as real increases. This is called a relativistic discriminator, and stabilizes training and improves the quality of samples.

While early GAN models required dataset sizes of up to tens of thousands of images, modern models can reach the same performance, or better, with smaller data sets. Karras et al. (2020) propose an adaptive discriminator augmentation mechanism that can stabilize training when training with limited data sets. This mechanism has since been adopted by different models with positive results.

Not much previous research investigating GANs in the domain of fabric pattern generation could be found. A previous study (Khalil et al., 2020) used Deep Convolutional GAN[4] (DCGAN) to generate textile patterns. However, this was the only variant of GANs tested and these results were not compared to other methods of image generation or other GANs. Next to this, the only objective metric used was the loss function values of the generator and discriminator over time. While these provide insight into the training process, they do not necessarily provide information about the quality of the images generated. They collected over 20 thousand images from various online sources and used this to train their model. While their results were somewhat promising, they were noticeably different from the state-of-the-art results GANs achieved with other datasets such as MNIST and CIFAR-10.

Figure 1. Final generated result from DCGAN. Adapted from Khalil et al.[16], Figure 8. As can be seen, samples created are noisy, and fail to resemble actual patterns or shapes.

DCGAN is an older model from 2015, and current models should be able to deliver better results. Their research concluded that their results could have been improved by a larger dataset, and more training time. The dataset used for this research will be composed of 47904 files. The larger size of this dataset could result in better performance with the same model.

4 Results

This chapter will lay out the results and also discuss the evaluation method used. Section 4.1 will discuss the evaluation method, followed by a showcase of the results of GANs in section 4.2. Section 4.3 will show the results of variational autoencoders.

4.1 Evaluation method

4.1.1 Fréchet Inception Distance

The Fréchet Inception Distance (FID)[11] is an improved version of the Inception Score, which is a metric said to correlate well with human judgment[5]. Both the Inception Score and FID score are popular metrics used for evaluating GANs. Both metrics attempt to score the GAN on how close in similarity images created by the GAN are to the original dataset. The lower the FID score, the more similar the images created are to the original dataset; a lower FID score indicates a better performing model. To accurately calculate the FID scores for the models used in this research, each model needs to generate a large number of images. An open-source Pytorch implementation created by Maximilian Seitzer[18] was used to calculate the FID scores.

4.2 Results GANs

Each model except StyleGAN2 was adapted from the open-source repository of Pytorch implementations of GAN models created by Erik Lindernoren[19]. Each of these models was tested on popular datasets and produced expected results. StyleGAN2 was adapted from another open-source repository which can be found at `https://github.com/lucidrains/stylegan2-pytorch`[17].

4.2.1 Parameters

During training every model except StyleGAN2 was trained with the following parameters:

- Image size: 64
- Number of epochs: 100
- Number of image channels: 3
- Learning rate: 0.0002

All other parameters were set to the default values, which can be found at Lindernoren's repository[19].

For StyleGAN2 the following parameters were used:

- Image size: 64
- Network capacity: 64
- Save every: 500 training steps
- Evaluate every: 1000 training steps
- Number of training steps: 100000

All other parameters were set to the default values, which can be found at the StyleGAN2 repository[17].

For the calculation of the FID score 32678 samples were created for each model.

4.2.2 GAN

GAN[1] was proposed by Goodfellow et al. in 2014, and since then GANs have become a popular method for image synthesis capable of state-of-the-art. However, the early models of GANs had issues such as mode collapse[1], training instability[3], and diminishing gradients[9]. Newer models attempt to remedy these issues.

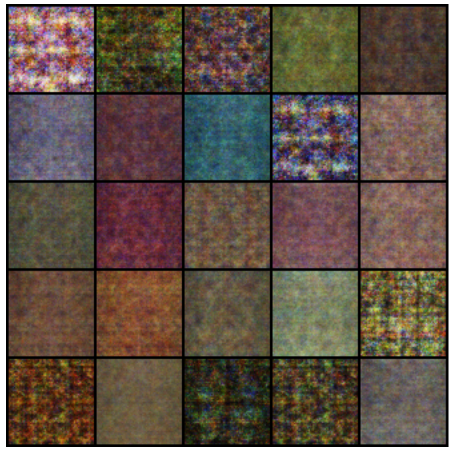

Figure 2. Final results from GAN after training. GAN fails to accurately replicate the images from the dataset, and mostly creates samples with only a background colour, but no actual patterns.

4.2.3 DCGAN

DCGAN[4] uses convolutional neural networks in GANs. This change helped stabilize the training of GANs and has since been adapted in many GAN models.

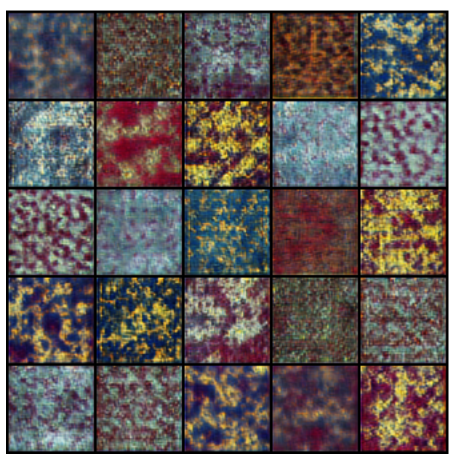

Figure 3. Final results from DCGAN after training. DCGAN performs better than GAN as the generated images are sharper but still fails to create refined shapes and lines.

4.2.4 ACGAN

Auxiliary Classifier GAN (ACGAN) is a variant of GAN employing label conditioning to improve global coherence of images. In ACGAN, the model is class conditional, but with an auxiliary decoder that is tasked with reconstructing class labels[10].

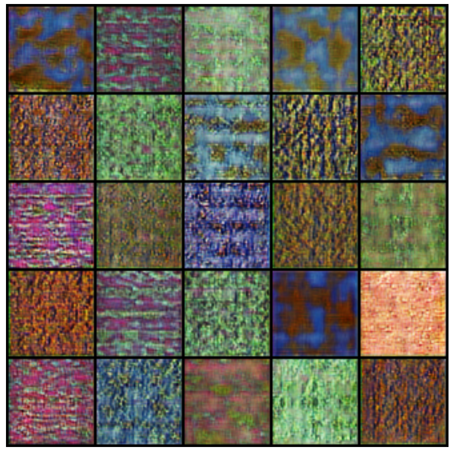

Figure 4. Final results from ACGAN after training. Similar to DCGAN, the resulting images are often fuzzy and do not resemble any real patterns. However, the relations between colours seems to more accurately resemble that of actual patterns as compared to DCGAN, there are fewer "spots" of colour and the borders between colours are more distinct.

4.2.5 WGAN Gradient Penalty

Wasserstein GAN (WGAN)[6] is a variant of GAN, which provides an alternative to traditional GAN training, by minimizing the Wasserstein distance. WGAN Gradient Penalty (WGAN-GP) is an alternative variant of WGAN that penalizes the norm of the gradient of the discriminator with respect to its input[8]. This change allows stable training with almost no hyperparameter tuning.

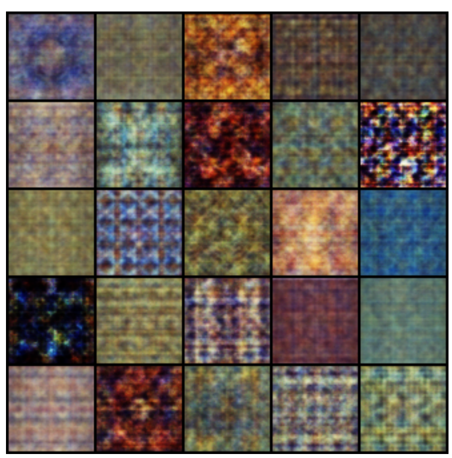

Figure 5. Final results from WGAN-GP after training. Different from DCGAN and ACGAN the presence of patterns seems to be much more clear, but these patterns are blurry.

4.2.6 Relativistic GAN

Relativistic GAN is a variant of GAN that uses a relativistic discriminator, which estimates the probability that a sample from the dataset is more realistic than a generated sample[12]. Relativistic Average GAN estimates the probability that a sample from the dataset is more realistic than a generated sample, on average.

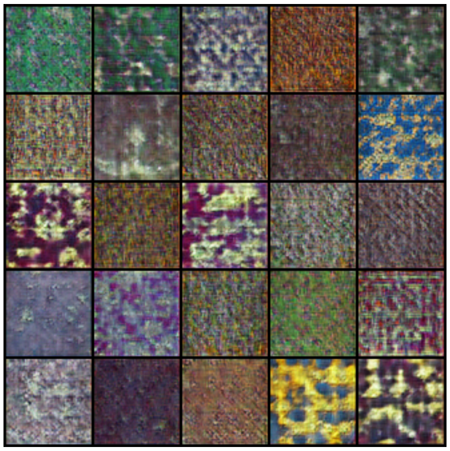

(a)

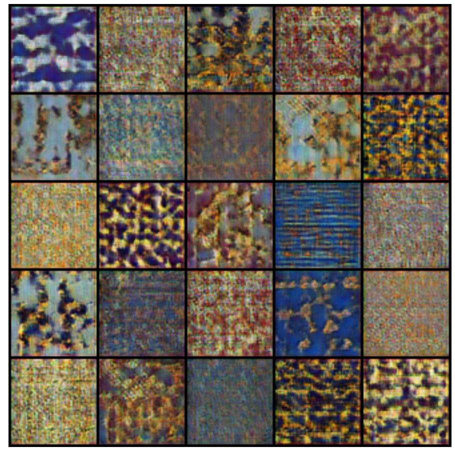

(b)

Figure 6. Final results from Relativistic GAN (a) and Relativistic Average GAN (b) after training. Compared to previous models, the images are sharp and some contain resemblances of simple patterns. However, in Relativistic Average GAN noise seems to be more apparent, resulting in fuzzy images.

4.2.7 StyleGAN2

StyleGAN2[14] is a model that further develops the performance of the style-based architecture called StyleGAN[13]. With changes to model architecture and training strategy StyleGAN2 improved the performance of StyleGAN and yields state-of-the-art results. StyleGAN2 also allows using a relativistic discriminator.

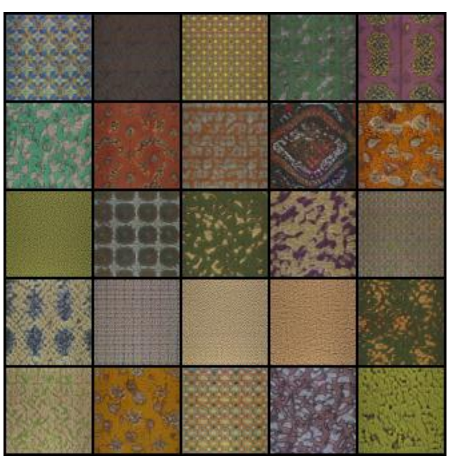

(a)

(b)

Figure 7. Final results from StyleGAN2 (a) and Relativistic StyleGAN2 (b) after training. Both models showcase the ability to create somewhat convincing samples, with clear improvement compared to the results of other GAN models shown in previous figures. Although both seem to be able to only recreate simple patterns, Relativistic StyleGAN2 does this more accurately.

4.2.8 Generated samples next to real samples

Figure 8. Eight randomly picked samples from publicly available fabric patterns on Vlisco's website[21] (top row) compared to one sample randomly picked from each model (bottom row) in order of GAN, DCGAN, ACGAN, WGAN (Gradient Penalty), Relativistic GAN, Relativistic Average GAN, StyleGAN2, Relativistic StyleGAN2. Samples from Vlisco's website were resized to the resolution of the generated samples. Improvement is visible, with the right-most samples showing much more resemblance to the patterns in the row above compared to the left-most samples.

4.2.9 FID scores of GANs

GAN Model	FID
GAN	297.5
DCGAN	202.1
ACGAN	157.5
WGAN-GP	231.5
Relativistic GAN	134.4
Relativistic Average GAN	165.8
StyleGAN2	69.0
Relativistic StyleGAN2	56.7

4.3 Results Variational Auto encoders

Figure 9. Eight samples generated by PixelSnail[7] using variational autoencoders. Adapted from work by Just Wallage[20]. The samples are fuzzy and do not resemble fabric patterns at all.

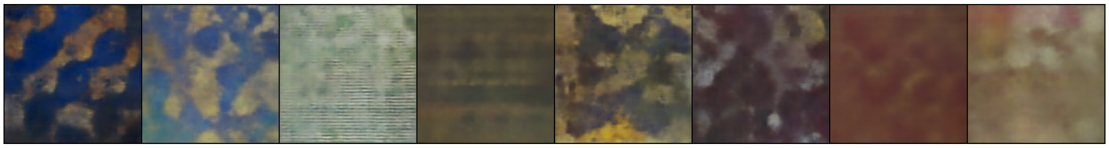

Figure 10. Eight generated samples by a variational autoencoder. Adapted from work by Just Wallage[20]. While these results are a clear improvement from those of PixelSnail (See Figure 9), the samples are blurry and lack refined shapes or symmetry.

5 Discussion

The results show clear improvements in the capability of GANs to create photo-realistic fabric pattern samples compared to the early models. The samples created by GAN were mostly just background colours with some noise. DCGAN shows some more variance, but the samples do not contain any smooth lines or actual shapes, which is a result similar to that of previous research[16]. ACGAN is similar to DCGAN, although the relations between colours of fabric patterns seem to be more accurately represented. WGAN-GP has the same issues as GAN. While some samples seem to contain a slight presence of patterns and symmetry, these are blurry. Relativistic GAN and Relativistic Average GAN seem to produce sharper samples than DCGAN and ACGAN. Compared to the other models, the images created by StyleGAN2 are much sharper and coherent but seem unable to replicate complex patterns.

The FID scores show that modern models significantly outperform older models. A result that stands out however is that of WGAN-GP, that despite being more modern than models such as DCGAN, had a worse FID score. A possible reason for this could be that the WGAN-GP model used does not use convolutional layers, with the only other model not using convolutional layers being GAN, the worst performing GAN model. This indicates that convolutional networks play a significant role in the creation of accurate samples for fabric patterns. The use of a relativistic discriminator has a positive effect on the results, which is especially evident in the results of StyleGAN2. Relativistic Average GAN performs worse than its counterpart. This could be due to the large variety of the dataset interfering with the use of averages.

The FID scores are higher than those for other datasets such as MNIST and CIFAR-10, with state-of-the-art GAN models achieving single-digit scores for these datasets. Poor choice of training parameters or the dataset not fitting well with the models used could be the cause of this. The large variety of colours and shapes present in the fabric pattern dataset could also have been too large in relation to the size of the dataset for the model to learn to create accurate samples.

Another possible limitation could be the image resolution of the dataset and the generated images. The images used for training had an average resolution of 349 by 468 pixels, with a standard deviation of 0.59 in width, and a standard deviation of 0.86 in height. Usually, all images in a dataset have the same resolution or aspect ratio. The images produced had a resolution of 64 by 64 pixels. This low resolution was chosen due to several reasons. As the resolution increases, the time needed for training can increase significantly, as the model will take more time per iteration and more iterations will be needed to reach the same quality of samples created. Next to this, the requirements for other resources such as memory will also increase. Due to limited resources and time, the priority was set on first evaluating the performance on GANs using these lower resolution samples. While it is expected that higher resolution samples will not deviate too much from lower resolution samples, the possibility exists that higher resolution samples will be able to capture the complexity of fabric patterns better and create better samples, but intermediate results for higher resolution models do not support this idea.

Similar to those of GANs, the results from variational autoencoders show the samples generated are still lacking compared to the performance on other domains. Compared to GANs, the results from variational autoencoders seem to be somewhat more blurry and lack structured figures and properties such as symmetry, which state-of-the-art GANs were able to replicate for simple patterns. The results from variational autoencoders are similar to those of DCGAN and the relativistic GAN models in terms of structure and figures although the variational autoencoders used seem to produce more refined samples as the images produced by the GAN models were much more fuzzy. The samples produced by PixelSnail lacked coherence and structure, performing worse than any GAN model.

After discussion with our contact at Vlisco, the general conclusion is that these results show that GANs cannot yet directly be applied to generating fabric patterns, as the results are not yet of sufficient quality. However, the insights these results provide are still of value and motivate future research on a larger scale. This would make this research a valuable first step of a larger process to reach the point where GANs could create patterns of similar quality as those currently produced by Vlisco.

6 Responsible research

This section will reflect on the possible ethical aspects of this research and discuss the steps taken to increase the reproducibility of the results. Lack of reproducibility is a common issue in artificial intelligence and machine learning research where the findings of a paper cannot be reproduced by another party. This could be because of several factors, which will be outlined in this section. Next to this, this section will discuss the possible influence of bias on the results of this research.

6.1 Ethical Implications

While this research did not necessarily handle any personal information, irresponsible conduct could still have implications. First of all, the data used for this research is sensitive and care should be taken to avoid leaking or sharing sensitive data. Due to this factor, the dataset used was stored privately and care was taken to avoid sharing the dataset through unsafe means. Next to this, the possible implications of this research should also be kept in mind. Designers can need multiple days or even longer to come up with a single design that meets all criteria, while an algorithm after training could generate thousands if not more of these designs per day. This could lead to a situation where human employees are replaced by a machine learning algorithm, which is a common concern regarding automation and machine learning. While current results do not suggest this happening soon for this domain, this should still be kept in mind.

6.2 Code and reproducibility

This research used open-source implementations of GAN models, and while these implementations were sometimes adapted to work with the relevant dataset or to aid in the measurements of results, the core of the actual models should remain intact, and these models should be able to be used by another party to produce similar results. The code repositories for every model used in this research can be found in the reference list, or by citation in the relevant section.

Reproducing these results accurately might be difficult. While the dataset used for this research cannot be shared, a similar dataset might become available and this could be used to verify or improve the findings of this research. But this is not the only possible issue. Because GANs are dependent on hyperparameters, the difference in results can vary extremely because of a small difference in these parameters. Because of this, the parameters used for the results shown for each model are shared in section 4.2.1.

Another issue with the reproducibility of machine learning research is due to the nature of machine learning; random numbers being used to start training and explore the search space for optimal solutions. Another factor could be the environment these models were trained in, with relation to the hardware used and possible software influences. These factors could lead to small deviations in the results.

6.3 Bias

Bias can play a significant role in research, both consciously and subconsciously. Because of this, it is important to be aware of the possible influences bias can have, and try to remedy these. In artificial intelligence and machine learning research the selection of data can have a large influence on the results. Because the dataset used for this was provided by Vlisco, and not handpicked by ourselves, we cannot know for sure whether or not certain decisions were made that would exclude or include certain samples, and what influences these decisions could have had on the results.

Another possibility could be in the selection of which results to present. While hundreds of thousands of images were generated for this research, only a handful was shown, which could lead to possible cherry-picking to give a false image of the capabilities of the models used. In an attempt to remedy this, three things were done. First, the images shown show multiple samples created in a 5 by 5 grid, and all of these images were created at the last iteration of training. Secondly, all other samples shown were picked randomly to avoid bias. Lastly, the FID score was used as an objective metric, for which a large number of images were generated and used to evaluate the model. Creating a large number of images avoids getting "lucky" samples and getting a much higher score than is realistic for the given model.

7 Conclusion and future work

7.1 Conclusion

This research aimed to investigate the performance of generative adversarial networks (GANs) on fabric pattern generation. While generative adversarial networks have been used in many domains with positive results, fabric pattern generation remained largely unexplored. Next to researching the performance of GANs on fabric pattern generation, these results were also compared to those of another popular method of image synthesis: variational autoencoders. The samples created by the GAN models used show clear development in the performance of GANs on fabric pattern generation. However, these results are not yet of the same quality as for other domains such as people's faces and other physical objects such as flowers. Convolutional layers seem to play an important role in generating sharp images, as models that did not use convolutional layers performed much worse than models that did, and created blurry samples. Due to limited resources and time, partly due to receiving the full dataset later than expected, higher resolution models could not yet be fully trained and evaluated, but current results for higher resolution models are similar to those of lower resolution samples, which were 64 by 64 pixels. Overall, the performance of GANs on fabric patterns is promising and has clearly improved over the years, but current results are not yet of the desired quality that GANs have been able to achieve for other domains. This means that while GANs cannot yet produce images of the same quality as human designers, they could already serve as a source of inspiration for designs. For Vlisco this research could function as a first step, motivating further research on a larger scale. Present findings also motivate further research into the influence of non-uniform datasets on the quality of images created by GANs.

7.2 Future work

Based on present findings, the only models capable of producing samples of somewhat decent quality are state-of-the-art models. Because of this, more research could be done into investigating the performance of a larger multitude of modern deep learning generative models, or future research investigating newer models and comparing the results to those of current models.

Another future research topic could be the generation of higher resolution fabric patterns, as the higher resolution samples would have more practical value and could give more insight into the current strengths and weaknesses of GANs with regards to this particular domain.

Research could be done into what qualities of a dataset are desirable to optimize the performance of GAN models, and how much differences in these qualities can affect performance. The dataset used contained images with aspect ratios that were not 1:1, which is uncommon for image datasets, next to a variance to the widths and heights of the images, which means there was a variance in aspect ratio as well, which could have lead to difficulties for the model to learn how to create square-sized samples that all have the same resolution. More research could be done into whether or not this had a negative impact on the quality of samples generated, and if this is the case how these effects can be remedied.

References

[1] Ian J. Goodfellow et al. *Generative Adversarial Networks*. 2014. arXiv: 1406.2661 [stat.ML].

[2] Diederik P Kingma and Max Welling. *Auto-Encoding Variational Bayes*. 2014. arXiv: 1312.6114 [stat.ML].

[3] Ian J. Goodfellow. *On distinguishability criteria for estimating generative models*. 2015. arXiv: 1412.6515 [stat.ML].

[4] Alec Radford, Luke Metz, and Soumith Chintala. *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*. 2016. arXiv: 1511.06434 [cs.LG].

[5] Tim Salimans et al. "Improved Techniques for Training GANs". In: *Advances in Neural Information Processing Systems*. Ed. by D. Lee et al. Vol. 29. Curran Associates, Inc., 2016. URL: https://proceedings.neurips.cc/paper/2016/file/8a3363abe792db2d8761d6403605aeb7-Paper.pdf.

[6] Martin Arjovsky, Soumith Chintala, and Léon Bottou. *Wasserstein GAN*. 2017. arXiv: 1701.07875 [stat.ML].

[7] Xi Chen et al. *PixelSNAIL: An Improved Autoregressive Generative Model*. 2017. arXiv: 1712.09763 [cs.LG].

[8] Ishaan Gulrajani et al. *Improved Training of Wasserstein GANs*. 2017. arXiv: 1704.00028 [cs.LG].

[9] Xudong Mao et al. *Least Squares Generative Adversarial Networks*. 2017. arXiv: 1611.04076 [cs.CV].

[10] Augustus Odena, Christopher Olah, and Jonathon Shlens. *Conditional Image Synthesis With Auxiliary Classifier GANs*. 2017. arXiv: 1610.09585 [stat.ML].

[11] Martin Heusel et al. *GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium*. 2018. arXiv: 1706.08500 [cs.LG].

[12] Alexia Jolicoeur-Martineau. "The relativistic discriminator: a key element missing from standard GAN". In: *CoRR* abs/1807.00734 (2018). arXiv: 1807.00734. URL: http://arxiv.org/abs/1807.00734.

[13] Tero Karras, Samuli Laine, and Timo Aila. "A Style-Based Generator Architecture for Generative Adversarial Networks". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019.

[14] Tero Karras et al. *Analyzing and Improving the Image Quality of StyleGAN*. 2020. arXiv: 1912.04958 [cs.CV].

[15] Tero Karras et al. *Training Generative Adversarial Networks with Limited Data*. 2020. arXiv: 2006.06676 [cs.CV].

[16] Saad Khalil et al. "Title: Generation of Textile Patterns Through Generative Adversarial Networks Generation of Textile Patterns Through Generative Adversarial Networks". In: Mar. 2020.

[17] Wang Phillip. *StyleGAN2 Repository*. https://github.com/lucidrains/stylegan2-pytorch. 2020.

[18] Maximilian Seitzer. *pytorch-fid: FID Score for PyTorch*. https://github.com/mseitzer/pytorch-fid. Version 0.1.1. Aug. 2020.

[19] Lindernoren Erik. *PyTorch GAN repository*. https://github.com/eriklindernoren/PyTorch-GAN. 2021.

[20] Just Wallage. *Graphic Design with Machine Learning: Fabric Pattern Generation using a Variational Autoencoder*. 2021.

[21] *Vlisco website*. https://www.vlisco.com/world-of-vlisco/design/fabric-stories/. Accessed on June 24th, 2021.