MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 34 Issue 2, July-Dec 2024, Pages: 497-507

OPTIMIZING GENERATIVE AI BY OVERCOMING STABILITY MODE COLLAPSE AND QUALITY CHALLENGES IN GANS AND VAES

Mariya Jaweed, BE in AI & DS Student, Stanley College of Engineering and Technology for Women
Raniya Farhan Shaikh, BE in CSE Student, Muffakham Jah College of Engineering and Technology

ABSTRACT

This paper reviews the current state of Generative AI, focusing on two of the most prominent models: GANs, an abbreviation for Generative Adversarial Networks and VAEs, an abbreviation for Variational Autoencoder. Despite the great potential of these models in producing high quality synthetic data, there are some issues that have been observed which include stability, mode collapse and quality of the generated results. There are many problems that arise during GAN training, including instability, which causes such problems as mode collapse, when the generator generates a small number of images or similar images. While running through the network, VAEs are observed to be more stable and less noisy than the GANs but they are not as accurate as the GANs in generation of samples. The mentioned limitations have been addressed by recent developments including Wasserstein GANs, feature matching, progressive GANs, and a combination of both such as VAE-GANs. These innovations are intended to increase stability and sample quality to employ new loss functions, training methods, as well as the new architectures that incorporate advantages of GANs and VAEs. However, several issues are still unresolved, these are computational cost, growth, and the question of the potential malicious use of the generative AI tools. This paper also discusses potential future research topics in the field, including self-supervised learning, combined multimodal methods, and the introduction of ethical measures when implementing generative models. The purpose of this review is to present the current state of the art, as well as current and potential issues in generative models, and different strategies to further increase the performance of these algorithms.

Keywords: Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Stability, Mode Collapse, Optimization, AI Quality, Deep Learning

1. INTRODUCTION

With the help of Generative AI, especially GANs and VAEs, there has been a great improvement in the generation of synthetic data of good quality. GANs, introduced by Ian Goodfellow in 2014, utilize a competitive framework between two neural networks: two major components, namely a generator and a discriminator. This adversarial setup has turned out to be useful for synthesis of realistic images, videos and other types of data. Whereas, VAEs learn a probabilistic model for the input data to obtain a continuous and smooth representation in the latent space for interpolation and data generation.

Both architectures have been successful in different uses such as image synthesis, data augmentation, and anomaly detection. But, here GANs and VAEs have some critical issues that downplay their effectiveness to a certain extent. Training GANs may be unstable where the generator and discriminator do not meet at an optimum or may produce deleterious results such as collapsing where the generator offers only a few kinds of solution. VAEs on the other hand, are more stable in terms of the latent space but the produced samples are blurry or of low quality due to the regularization of the latent space at the cost of fine details reconstruction. These issues reduce their ability to produce high fidelity information, especially in those areas where refined detail is required. This paper seeks to discuss these main challenges, stability, mode collapse, and quality, and discuss the different optimization methods that have been developed in the current literature to solve them. This review helps to understand how optimization strategies improve the GANs and VAEs to make them suitable for real-world applications by outlining the most critical problems and solutions. The paper will not dwell on peripheral issues like computational efficiency of the methods used to solve these problems, but will instead concentrate on the technical sides of these challenges.

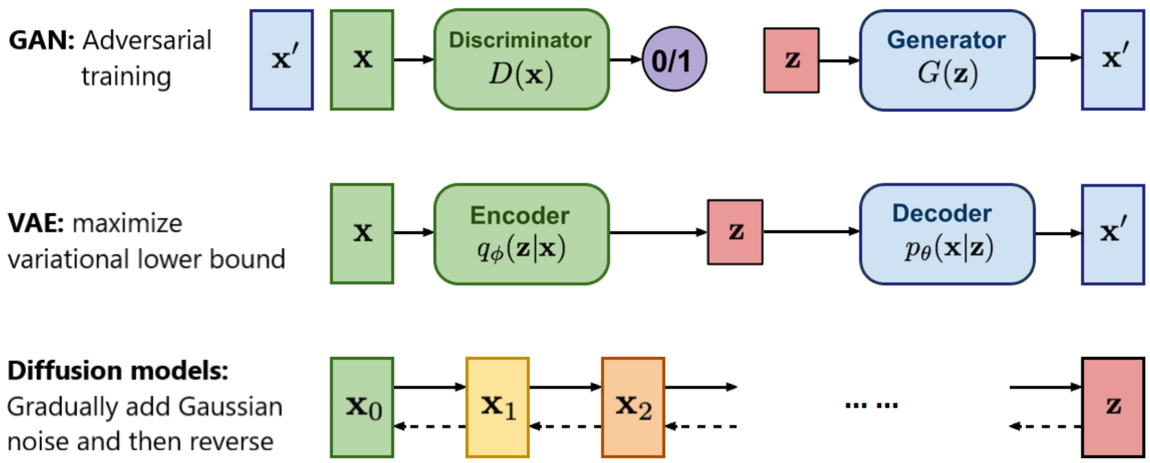

Fig No. 1 GANs and VAEs Collapse

Background

The field of generative AI has experienced great progress in the last ten years and has spread across various areas from computer vision to natural language processing. Based on the generative models, the most prominent are GANs and VAEs. As mentioned by Ian Goodfellow in 2014, GANs consist of two neural nets, the generator and the discriminator where the two are in a game-like environment that produces realistic synthetic data. VAEs are fundamentally different from VAEs as they are probabilistic models that learn a set of factors which could generate new data points. Both models have proved to be very successful in producing images, audio, videos and even texts. But they also have their drawbacks. In particular, the stability of GANs during the training process is a critical problem since the generator and discriminator may become unbalanced, and the generator cannot generate realistic data even if the discriminator works well. Another famous issue is mode collapse, when the generator provides a finite number of outputs instead of the variety of the data distribution. VAEs, which are less likely to collapse, suffer from reconstruction loss, and the produced samples may be blurry or of lower resolution compared to GANs. These issues need to be addressed in order to enhance generative models' quality and their relevance for practical use cases.

Objective/Purpose

The objective of this paper is to discuss the primary problems associated with GANs and VAEs concerning stability, mode collapse, and quality. We will also look at the optimization techniques that have been advanced in recent literature to solve these problems. With respect to the various approaches that have been proposed in the literature to enhance the performance of these generative models, this paper presents a survey of the current literature. The purpose of this paper is not only to highlight what aspects need to be improved, but also to outline the direction for further research aimed at making generative AI models even more reliable.

Overview of Structure

The paper is structured as follows: we will then give an overview of the historical background and current status of GANs and VAEs and discuss the evolution of the models and the research on their drawbacks will focus on the technical details of the problems in question, outlining the issues of stability, mode collapse and quality and then present the recent developments and methods designed to counter these issues. After that, we will proceed to the discussion of the rest of the current issues. Lastly, we will discuss future research implication and conclude this paper with the findings.

2. LITERATURE REVIEW

In Generative AI there is much progress in the recent years especially in Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), which make it possible to generate synthetic data of high quality in many sectors. Both GANs and VAEs, but are still have their own issues, difficulties in terms of the stability, mode collapse, and the quality of the generated samples. GANs work by using a generator and a discriminator where the generator's aim is to generate data which is as real as possible while the discriminator aims at distinguishing between the real data and the generated data. This setup often results in instability during training, and one of the common problems which arise is mode collapse, in which the generator creates a small variety of outputs, not covering the whole spectrum of the actual data. There has been recent work done on these issues with the advent of inception GAN, Wasserstein GAN, etc that use different loss function based on the Wasserstein distance to get more stable gradients and avoid mode collapse. Further, gradient penalties, spectral normalization, and feature matching have been suggested to enhance the stability and data diversity of the generated data.

On the other hand, VAEs use a probabilistic approach to modeling the underlying manifold in which data is generated from. Although this approach is more effective in terms of stability during training, it is reproached for generating blurry images since the model seeks to obtain a smooth latent space at the cost of sharpness. Another issue of the VAEs is that the diversity of the points in the latent space can be inversely proportional to the quality of the generated outputs. Despite the improved disentanglement, which in turn results in better interpretable latent representations, this approach sacrifices the sharpness and realism of the generated data. These problems have been addressed in the recent developments including perceptual loss and attention mechanisms that help to improve the quality of the output. The application of the GANs and VAEs for image generation also has attracted the researchers' attention recently and the investigation of the composite models, which integrate the advantages of the GANs and the VAEs, has begun.

Another such model is the VAE-GAN, which combines the stochastic element of VAEs with the adversarial training of GANs in order to generate better samples. These models try to give a middle ground between the continuous data generation and the interpretable nature of VAEs and the fine and detailed samples from GANs. Another type of the mentioned hybrid model is the adversarial auto encoder that uses adversarial training in the structure of autoencoder, and thus provides better representation learning and more accurate data generation. Several new developments have been made in the last few years to overcome both stability and quality problems in generative models. For instance, Progressive GANs slowly expand the network during training at an early stage, the model learns to generate low resolution data and later progresses to high resolution outputs, this approach enhances both stability and quality of the samples. Also, the architecture of style-based GAN provides more options in terms of style and detail of the generated data, and reduces distortion.

Nevertheless, several obstacles still exist in the field. GANs remain challenged in managing the balance between the generator and discriminator, and stability in training may minimize the variance in generated samples. For VAEs, the competition between the quality of reconstructions and the ability to achieve disentanglement in the latent space is still a challenge. In addition, both GANs and VAEs often encounter computational challenges by virtue of the optimization processes and models' growth, which would lead to concerns of scalability and computational costs. Implementation of generative models also has challenges on the practical level. Even with enhanced model quality and stability, these models have a problem of domain shift, where they do not perform well when applied to new data environments. They are also prone to adversarial attacks, in which their performance tends to be compromised. Some of these challenges include the capability of the model in creating deepfakes or fake data for unethical use which have drawn researchers' attention.

Coordinating responsible and ethical use of generative AI is necessary as these models are applied in production environment. Summing up, one can note that the current research provides a certain number of advancements in overcoming the limitations of GANs and VAEs, although there are still certain trends that require further attention: computational efficiency, model robustness, and ethical issues. The future of generative AI will be defined by further improvements of optimization algorithms, integration of new hybrid systems and by addressing practical issues of deployment that will help to produce high quality synthetic data for a range of applications.

3. TECHNICAL CONCEPTS AND METHODOLOGIES

GANs (Generative Adversarial Networks)

GANs are composed of two neural networks: the generator, the discriminator. The generator generates new data and the discriminator analyses whether the data is genuine or a fake sample. The two networks are trained in a minimax game: the generator is trying to deceive the discriminator and on the other hand the discriminator is trying to get the data right. This adversarial process encourages the generator to improve continuously, resulting in highly realistic generated samples. Despite their success, GANs face stability issues during training. The balance between the generator and discriminator is critical—if one outpaces the other, training may become unstable. For instance, the vanishing gradient problem occurs when the discriminator becomes too powerful, causing the generator to receive no meaningful gradients and halting learning. Conversely**,** exploding gradients can lead to excessively large updates, destabilizing

MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 34 Issue 2, July-Dec 2024, Pages: 497-507

training. Additionally, mode collapse is a significant challenge in GANs, where the generator produces a limited variety of outputs, often creating a few "safe" samples that consistently fool the discriminator.

This reduces the diversity of generated data, an issue particularly detrimental for applications requiring varied outputs.

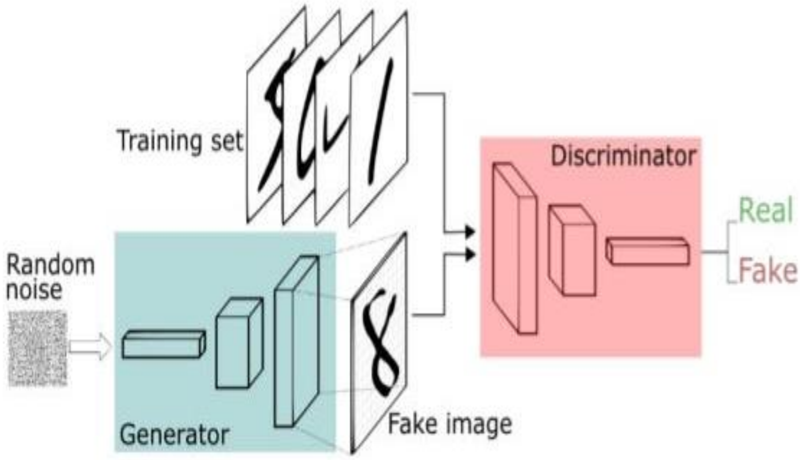

Fig no 1 GANs (Generative Adversarial Networks)

VAEs (Variational Auto encoders)

VAEs differ fundamentally from GANs in their approach to data generation. A VAE encodes input data into a latent space through an encoder, and then decodes it back into the original space. The process optimizes the reconstruction **loss** (how accurately the decoded data matches the original) and the KL divergence, which regularizes the latent space to be close to a known distribution (typically a Gaussian). This approach offers stability during training but often results in blurry or low-quality samples, as VAEs tend to smooth over fine details to maintain a more continuous latent space. The main challenge with VAEs lies in sample quality, as the reconstruction loss and the constraints placed on the latent space often limit their ability to generate sharp, detailed images. While VAEs are robust and stable, they frequently generate images that lack the realism and detail seen in GAN-generated data.

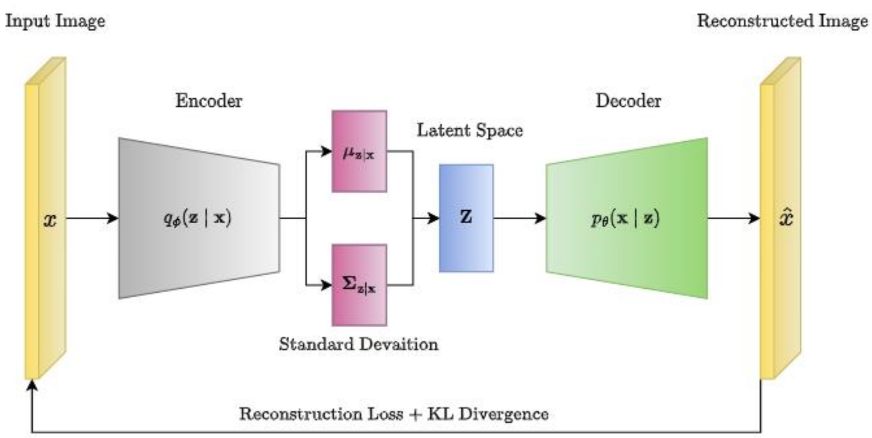

Fig no 2 VAEs (Variational Autoencoders)

Common Challenges

Both GANs and VAEs share common challenges that impact the quality of generated data. In GANs, the instability of training (due to vanishing/exploding gradients and mode collapse) often leads to poor convergence or limited diversity. For VAEs, the trade-off between latent space regularization and detailed data reconstruction results in blurry or low-resolution outputs.

Optimization Approaches

Recent advancements have sought to address these challenges. For GANs, there are WGANs that use the Earth Mover's distance as a more stable measure of the distance between two probability distributions and, hence, the gradients do not vanish or explode. Spectral normalization has also been proposed to stabilize the training of both the generator and discriminator using weight scaling in the network, which increases training instability. Feature matching and minibatch discrimination are among the approaches which are used to prevent mode collapse and maintain the variety of the produced samples. In the case of VAEs, disentangled representations have been suggested in order to enhance the sample complexity and quality of the synthesized samples. This technique makes it possible for each of the latent variables in the VAE to be a unique factor of variation, thus producing a variety of detailed outputs. Moreover, incorporating perceptual loss functions and adversarial training into hybrid models (VAE-GAN) can assist VAEs in generating better better and more photorealistic images than VAEs and GANs. To sum up, both GANs and VAEs represent the breakthrough in generative AI, but they currently have many issues connected with stability, diversity, and sample quality.

These problems and others are tackled by optimizations such as WGAN, spectral normalization and disentangled latent representations and are still improving the generative models.

4. RECENT TRENDS AND ADVANCES

Improving Stability in GANs

Another critical problem of GANs is the instability of the training phase, though the recent improvements made in this area have brought some relief. WGANs which were proposed by Arjovsky et al., replace the objective function of the typical GANs, the Jensen-Shannon divergence, with the Earth Mover's distance or Wasserstein distance. This new loss function gives rise to smoother gradients especially when the generator and the discriminator are far from being synchronized thus improving the training stability.

The gradient penalty introduced in WGAN-GP enhances the stability by enforcing the Lipschitz continuity constraint to make the discriminator gradients well-conditioned. Another effective method for making GANs stable is batch normalization where the activations within each mini-batch are normalized. This minimizes internal covariate shift and lets the generator work better, enhancing both convergence and performance. Also, adaptive optimization methods like RMSProp and Adam have been applied in this work to adapt the learning rates to avoid problems like gradient explosion or vanishing gradients.

Tackling Mode Collapse

One of the main problems of GANs is mode collapse, in which the generator outputs only a small number of modes and does not represent the range of data distribution. To deal with this problem, several techniques have been devised. Feature matching is another approach that aligns

MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 34 Issue 2, July-Dec 2024, Pages: 497-507

the generator to generate outputs that have properties of the middle layers of the discriminator thus avoiding the generation of similar samples. Minibatch discrimination increases diversity because the discriminator can see multiple samples in a minibatch at a time and, therefore, is less likely to be deceived by the generator with only a few modes. Unrolled GANs the concept of adversarial training by thinking a few steps further in the optimization process to minimize the chances of the generator outsmarting the discriminator. The aims of multi-objective GANs, which are included in the training process, are such objectives as diversity and or novelty, which make the generator search for more outputs and minimize mode collapse.

Enhancing Quality in GANs and VAEs:

More recent work has been geared on enhancing the quality of the data that are generated from GANs and VAEs.

For example, progressive GAN training is to increase the network complexity during the training process, from low to high resolution. This strategy aids the model to converge and enhanced the quality of generated samples, especially when the images are of high resolution. Attention mechanisms, popular in natural language processing and computer vision in recent years, are applied to GANs in order to help the generator to focus on a specific area of an image and produce more detailed and realistic fake data. As for loss functions, perceptual loss has emerged recently, where the similarity between images is not evaluated by pixel difference, but through comparing feature representations of these images extracted by pre-trained networks such as VGG. This assists in retaining the small features and architecture in the images thus enhancing the quality of the generated images. More so, style-based GANs like StyleGAN present a new architecture where the generator's latent space is divided in different layers for different styles and produces very realistic images.

Hybrid Approaches

To extend the idea of both the variability and the quality of the synthesized samples, the latest studies have focused on the integration of GANs with VAEs. One such hybrid is the VAE-GAN which is a combination of the probabilistic model of the VAEs and the adversarial training of the GANs. This approach utilizes the fact that VAE has a stable latent space for sample generation and GAN for high quality realistic samples. The GAN part further improves the quality of the generated data while the VAE part ensures data is placed in a well organized latent space hence high diversity. Another fine hybrid approach combines CVAEs with GANs since they enable the generation of conditional data (for example, images conditioned on labels or text) with both high level of variability and high level of details. These kinds of models suggest that there is a potential for obtaining the best of both VAEs and GANs, which results in improved performance in terms of sample diversity, stability, and quality. They show that there is still active work on addressing the issues with GANs and VAEs, and thus, there is constant progress in the generative AI field to create a better generative model that can generate a diverse and high-quality synthetic data.

5. CURRENT ISSUES AND LIMITATIONS

In GANs

Even though, there has been much improvement on GANs, the following remain as issues of concern especially the balance between the generator and the discriminator. What is more is that GANs are infamous for this fine balance between these two elements.

If the generator becomes stronger than the discriminator, the problem occurs in the diversity of the generated samples (mode collapse); if the discriminator is too strong, gradients vanish and learning slows down and the generator fails. This imbalance can still cause training instabilities and, while tricks such as WGANs or gradient penalties can help with that, they are not a cure-all. Furthermore, the key limitation of GANs is that they are typically unstable and have low variability at the same time. Some of the strategies that makes training stable (for example the gradient penalty in WGAN-GP) can help with convergence but might limit the range of variation in the data generated by the model, thereby providing a tradeoff between generating realistic data and generating a diverse set of variations of the data.

In VAEs

While VAEs are less sensitive to the problem than GANs, they still face the choice between improving sample quality and achieving better disentanglement in the latent space. The goal of VAEs is to regulate the latent space by means of KL divergence and to obtain a meaningful latent space. However, this regularization is usually accompanied by the loss of the reconstruction quality, as the model aims at reconstructing the big picture of the data rather than details. Consequently, generated samples may appear blurry or can have low fidelity in the details, which is especially the case with realistic image synthesis. In addition, disentanglement of the latent space – the capability to guarantee that each latent variable corresponds to a different source of variation – is still an issue. Beta-VAE and factorized VAEs are the recent approaches that try to solve this problem, but the problem of full disentanglement while preserving the high quality of generated reconstructions has not been solved yet.

Optimization Complexity

The improvement of both GANs and VAEs is usually associated with high computational costs. WGANs and perceptual loss functions enhance the quality and trainability of generative models but at the same time, they are computationally expensive. For instance, WGAN-GP entails computation of the gradient penalty, while with the perceptual loss functions one has to pass the data through a pre-trained network, which is time consuming. For training GANs and VAEs, as models get more complex, the need for better and faster hardware, and longer training times makes it unimplementable on a large-scale real-world problem. Moreover, these methods can be also faced with significant problems in terms of time and space complexity, especially when applied to very large datasets, which are not suitable for real time or large-scale generation such as video generation or interactive environments.

Practical Deployment Challenges

Despite the recent progress, generative models have several issues when it comes to deployment. The first and most obvious issue is scalability. Although GANs and VAEs perform well in small scale problems, the extension of these models to generate high resolution data or large data sets is computationally expensive and memory intensive. Domain adaptation is another problem; generative models trained in one dataset perform poorly when applied to other domains or conditions. For instance, a model trained on a dataset of faces will perform poorly when applied to medical imaging as the distribution of data is quite different. Robustness is also an issue: generative models are normally affected by input noise and may generate artefacts or unrealistic outputs when tested with variations or adversarial examples. These challenges limit the use of generative models in the important areas such as healthcare, autonomous vehicles, and content generation where model robustness and flexibility are paramount. Therefore, GANs and VAEs have been advanced significantly in recent years and still, they have some essential issues that include training stability, diversity, reconstruction quality, computational complexity, and the problem of real-world application. Solving these problems will be crucial for further development of generative models, their improvements in terms of stability and scalability, as well as the expansion of their domain of usage.

6. FUTURE DIRECTIONS

The future trends for generative AI, especially for GANs and VAEs are aimed at improving the results of the models, at broadening the areas of application, and at solving the ethical problems related to these technologies. One interesting avenue is the interaction of self-supervised tasks with generative models. Current self-supervised learning methods can be thought of as having the ability to effectively utilize large amounts of unlabelled data, increase the quality of the generated samples and build more reliable models while not necessarily requiring large amounts of labelled training data. This could potentially let generative models learn more complex and high-dimensional data distribution, and at the same time require much less labelled data. One of the new topics that need further investigation is the usage of energy-based models (EBMs) to improve the stability and quality of generative models. EBMs are proposed to minimize an energy function with respect to the data distribution, which can be a potential solution to the GAN training issue. Such models can afford finer-tuning control with the generated output and result in a better convergence during the training. Integrating EBMs with GANs or VAEs could give a better and less prone to training artifacts framework for synthesizing high quality data. Another important strand that will be pursued further is the development of hybrid and multimodal designs.

To achieve the improved performance, people incorporate the advantages of both GANs and VAEs and incorporate other generative approaches, including normalizing flows and autoregressive models. For instance, VAE-GAN hybrids have been designed, which has demonstrated potential for enhancing the quality and the amount of varied data produced. In future works, the hybrid models might build upon this by adopting ideas from other domains like attention mechanisms, transformer architectures, and disentangled representation learning to increase model's complexity and overall effectiveness. Last but not least, the ethical implications will be the primary driver for the future advancement of generative AI. With the advance in generative models the possibilities for malicious use such as deep fakes or any other synthetic data manipulation will also rise. Subsequent research will have to investigate methods for reducing bias and explaining, as well as being accountable for, the output of generative models. Further, structural approaches to identify and prevent the generation of the toxic or fake content will be inevitable in order to deploy these technologies responsibly.

7. CONCLUSION

There has been a lot of progress in generating models, especially GANs and VAEs over the past few years, and the application of these models is not limited to image synthesis but also data augmentation and creative content generation. However, problems such as instability of the training process, mode collapse, and sample quality have not been solved yet. Some of the recent development include WGANs, progressive GAN training, and hybrid VAE-GAN models which has improved on the stability, diversity, and realism of the models but it is apparent that the search

MSW MANAGEMENT -Multidisciplinary, Scientific Work and Management Journal
ISSN: 1053-7899
Vol. 34 Issue 2, July-Dec 2024, Pages: 497-507

for better Generative models that are robust, scalable, and efficient is still ongoing. As for the future trends, it is possible to expect that the work will continue with the help of more sophisticated methods, including self-supervised learning, energy-based models, and adversarial training with reinforcement learning as the potential ways of the model's improvement and flexibility. Furthermore, the mode collapse issue will be resolved through dynamic architecture, and the training algorithms will be optimized, which will be important to guarantee the quality of the produced outputs. However, the application of GANs and more generally generative models in hybrid and multimodal architectures, as well as the integration of different generative techniques, suggests the possibility of new future developments and new fields of use. However, autonomy raises important ethical concerns including how to avoid bias, prevent fake data generation and address misuse, all of which will be key when seeking to reap on the benefits of generative AI.

While the field evolves, there will be a need to ensure that the work done is both technically creative and ethically sound as the full potential of the generative models is realized. Thus, it can be concluded that with the current development, further work will be needed to eliminate the shortcomings of GANs and VAEs. The next generation of generative models shall look at both the technical and ethical aspects and shall be more efficient, reliable, and socially appropriate for the various application domains in Science, industry and Society.

REFERENCES

1. Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv Preprint arXiv:1411.1784*. https://arxiv.org/abs/1411.1784
2. Perarnau, G., van de Weijer, J., Raducanu, B., & Álvarez, J. M. (2016). Invertible conditional GANs for image editing. *arXiv Preprint arXiv:1611.06355*. https://arxiv.org/abs/1611.06355
3. Jaiswal, A., AbdAlmageed, W., Wu, Y., & Natarajan, P. (2019). Bidirectional conditional generative adversarial networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11363 LNCS, 216–232.
4. Máttyus, G., & Urtasun, R. (2018). Matching adversarial networks. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 8024–8032.
5. Kwak, H., & Zhang, B.-T. (2016). Ways of conditioning generative adversarial networks. *arXiv Preprint arXiv:1611.01455*. https://arxiv.org/abs/1611.01455
6. Odena, A., Olah, C., & Shlens, J. (2017). Conditional image synthesis with auxiliary classifier GANs. *34th International Conference on Machine Learning (ICML)*, 6, 4043–4055.
7. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). Improved techniques for training GANs. *Advances in Neural Information Processing Systems*, 2234–2242.
8. Wang, Z., Simoncelli, E. P., & Bovik, A. C. (2003). Multi-scale structural similarity for image quality assessment. *Conference Record of the Asilomar Conference on Signals, Systems, and Computers*, 2, 1398–1402.
9. Lee, M., & Seok, J. (2019). Controllable generative adversarial network. *IEEE Access*, 7, 28158–28169.
10. Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *32nd International Conference on Machine Learning (ICML)*, 1, 448–456.

11. Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). Layer normalization.
12. Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. *6th International Conference on Learning Representations (ICLR)*.
13. Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. *4th International Conference on Learning Representations (ICLR)*.
14. Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2018). Progressive growing of GANs for improved quality, stability, and variation. *6th International Conference on Learning Representations (ICLR)*.
15. Lin, Z., Fanti, G., Khetan, A., & Oh, S. (2018). PacGan: The power of two samples in generative adversarial networks. *Advances in Neural Information Processing Systems*, 1498–1507.
16. Im, D. J., Kim, C. D., Jiang, H., & Memisevic, R. (2016). Generating images with recurrent adversarial networks. *arXiv Preprint arXiv:1602.05110*. https://arxiv.org/abs/1602.05110
17. Saatchi, Y., & Wilson, A. G. (2017). Bayesian GAN. *Advances in Neural Information Processing Systems*, 3623–3632.
18. Jaiswal, A., AbdAlmageed, W., Wu, Y., & Natarajan, P. (2019). CapsuleGAN: Generative adversarial capsule network. *Lecture Notes in Computer Science (LNCS)*, 11131, 526–535.
19. Zhang, A., Han, X., Goodfellow, I., Metaxas, D., & Odena, A. (2018). Self-attention generative adversarial networks. *arXiv Preprint arXiv:1805.08318*. https://arxiv.org/abs/1805.08318
20. Durugkar, I., Gemp, I., & Mahadevan, S. (2017). Generative multi-adversarial networks. *5th International Conference on Learning Representations (ICLR) - Conference Track Proceedings*.
21. Nguyen, T. D., Le, T., Vu, H., & Phung, D. (2017). Dual discriminator generative adversarial nets. *Advances in Neural Information Processing Systems*, 2671–2681.
22. Chavdarova, T., & Fleuret, F. (2018). SGAN: An alternative training of generative adversarial networks. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 9407–9415.
23. Van den Oord, A., Kalchbrenner, N., & Kavukcuoglu, K. (2016). Pixel recurrent neural networks. *33rd International Conference on Machine Learning (ICML)*, 4, 2611–2620.
24. Van den Oord, A., Kalchbrenner, N., Vinyals, O., Espeholt, L., Graves, A., & Kavukcuoglu, K. (2016). Conditional image generation with PixelCNN decoders. *Advances in Neural Information Processing Systems*, 4797–4805.
25. Che, T., Li, Y., Jacob, A. P., Bengio, Y., & Li, W. (2019). Mode regularized generative adversarial networks. *5th International Conference on Learning Representations (ICLR)*.