The Journal of Computational Science and Engineering (TJCSE)
ISSN 2583-9055 (Media Online)
Vol 3, No 7, July 2025
PP 112–129

GANs and Beyond: Understanding Image Style Transfer with Deep Learning

D. Saketh Varma[1], D. Siva Sunil Prasad[2], V. Abhihith Naidu[3] and A.S. Venkata Praneel[4]

[1,2,3,4]Department of Computer Science and Engineering,
GST, GITAM (Deemed to be University), Visakhapatnam, AP, India.
sakethvarma2277@gmail.com[1], dvsivasunil2003@gmail.com[2],
vennelaabhihithnaidu@gmail.com[3] , pankired@gitam.edu[4]
Corresponding Author *: pankired@gitam.edu

Abstract

Image style transfer is a powerful deep learning method that combines one image's content and another's style. The multi-disciplinary method of style transfer has applications in digital painting, video restoration, medical imaging, architecture, and data augmentation. Non-photorealistic rendering and texture synthesis were the precursor methods to style transfer. But deep learning, especially Generative Adversarial Networks (GANs) and Convolutional Neural Networks (CNNs), has reached incredible progress within the field. Of the current methods, CycleGAN has disrupted style transfer because it can handle unpaired image-to-image translation, which proves to be incredibly beneficial for application in artistic rendering and medical imaging, where paired sets of data don't exist. In parallel, StyleGAN has brought in photorealistic synthesis with a high degree of control over the features of an image through the manipulation of latent space. Other attention-based models like StyTr2 and Adaptive Instance Normalization (AdaIN) enhance style transfer by allowing selective attention on areas of the image that matter. These innovations are useful in real-world applications across various fields. In medical imaging, CycleGAN normalizes data sets by eliminating color and texture variations to enhance diagnostic accuracy. StyleGAN has been helpful in data augmentation by creating synthetic yet realistic images that contribute to model training in applications such as facial recognition and content generation. Attention-based models are also improved by not letting style transfer skew important structural information, and as such, they are applicable in high-precision areas such as radiology and histopathology. Despite all this, many challenges are involved, such as high computational cost, training complexity, and limited scalability. Using many datasets, the following study must operate efficiently, minimize resource usage, and increase the style transfer model's adaptability. This current work is a thorough overview of the methodology already employed and its weaknesses, limitations, and applicability.

By rectifying these issues, style transfer models may be enhanced to maximize creativity further, improve automation, and strengthen the development of computer vision technology applications.

Keywords: CNN, deep learning, GAN, Cycle GAN, StyleGAN, StyTr2, Adaptive Instance Normalization (AdaIN).

1. Introduction

Image style transfer is a computer vision technique where one image's content is combined with another's style. The new image retains the structural details of the content image but has the artistic or visual style of the style image. It has become an effective tool in creative and practical applications, combining creativity with advanced deep-learning techniques. The process may convert a photograph into a painting in the style of Van Gogh or Picasso, with all shapes and compositions preserved from the original photo.

Image style transfer consists of two major components, namely, content and style. The content represents every structural and semantic detail of the image that expresses the actual idea of "what" by locating things and telling more about their shapes, spatial layouts, or arrangements of the items, like buildings, streets, and trees, in a capture of a city. Style is then expressed in all visual characteristics typical of an image by answering the question of "how"-the colors, textures, brushstrokes, and patterns, such as those representatives of Van Gogh's swirls or Picasso's geometric lines. Image style transfer allows for blending one image's structural content with another's artistic style, generating visually pleasing outputs. This has revolutionized digital art creation by automating the process of applying artistic styles and work using CNNs (Style Transfer via Gram Matrix) [1] to create stunning artwork.

Style transfer has moved beyond art, finding uses in medicine, architecture, and data augmentation. Medical imaging is linked with CycleGAN, which removes color variability in the images used for diagnosis and improves performance. Architects use style transfer to visualize period structures in modern attire, helping design exploration.

1.1 Technical Advancements and Research

Style transfer has driven advancements in deep learning, with approaches such as GANs for high-quality outputs (StyleGAN). Like those in StyTr2, attention mechanisms improve results by focusing on highly relevant image regions [2].

1.1.1 Types of Image Style Transfer

A deeper comparison of image style transfer methods in terms of computational efficiency, scalability, and performance across various domains highlights the strengths and limitations of each approach. Image-iteration-based methods, such as Neural Style Transfer (NST) and optimization-based techniques, rely on iterative gradient updates to refine image styles. While they produce high-quality results, they are computationally expensive and unsuitable for real-time applications. Additionally, their scalability is poor, as each new style requires re-optimization, making them inefficient for large-scale implementations. These methods perform well in artistic style transfer but struggle with handling complex textures and practical applications like medical imaging, where precision and efficiency are critical.

To address these limitations, model-iteration-based methods, such as Fast Neural Style Transfer, leverage pre-trained networks to enable real-time style application. This significantly improves computational efficiency by reducing the need for iterative updates. However, scalability remains challenging, as each style requires a separately trained model, increasing storage and training costs. These methods are particularly effective for real-time applications, such as augmented reality (AR) filters and mobile-based artistic rendering. Despite their speed advantage, they lack adaptability across diverse domains, limiting their use in fields that require flexibility in style application.

GAN-based methods, particularly CycleGAN and StyleGAN, offer a significant advancement by utilizing adversarial learning instead of iterative optimization. This reduces computational intensity compared to NST while generating highly realistic style transfers. However, GANs are still resource-intensive, requiring large datasets and powerful GPUs for practical training. Scalability is a strong advantage of GAN-based models, especially CycleGAN, which does not require paired datasets, making it ideal for applications in medical imaging, domain adaptation, and artistic transformations. On the other hand, StyleGAN excels in photorealistic image synthesis, character design, and artistic rendering, offering fine control over style manipulation. However, it requires substantial computational power and struggles with small datasets, leading to overfitting in specific scenarios.

More recently, attention-based models, such as StyTr2 and Adaptive Instance Normalization, have improved the efficiency and accuracy of style transfer. These models optimize computational efficiency by selectively focusing on important regions of an image, reducing unnecessary processing. However, including attention layers adds computational overhead, which may impact real-time applications. Unlike traditional model-iteration-based methods, attention-based approaches offer greater scalability, as they can adapt to multiple styles without requiring model retraining. These models have shown strong performance in medical imaging (e.g., MRI, histopathology), artistic rendering, and AR applications. Still, they often require fine-tuning for domain-specific datasets to achieve optimal results.

Figure 1: On the left, a photograph of a modern cityscape is presented, while the middle features the iconic painting "The Starry Night." The right side displays the output of applying Van Gogh's artistic style to the cityscape, resulting in a unique image that retains the original structure of the photograph while adopting the vibrant, swirling patterns characteristic of the painting. [3].

Before GANs, image style transfer relied heavily on optimization-based methods and CNNs to combine content and style. The key approach was image-iteration-based style transfer, where a noise image was iteratively optimized to align with the content features of one image and the style features of another.

1.1.2 CNN-Based Feature Extraction

- Early techniques used pre-trained CNNs like the VGG network to extract hierarchical features from images.
- The content features were derived from higher layers of the network, representing the semantic structure of the content image.
- Style features are captured using the Gram matrix, which measures the relationship between image styles. [4].

1.1.3 Loss Function Optimization

- The process optimized an image to minimize a total loss function, a weighted sum of content loss and style loss.
- Content Loss: The difference between the content features of the resulting image and the different graphic elements.
- Style Loss: The difference between the pattern features of the resulting image and the image style, measured using the Gram matrix.
- The optimization typically required multiple iterations, making the process computationally expensive. [4]

1.1.4 Limitations of Pre-GAN Methods

- Computational Intensity: The iterative process was slow and resource-intensive.
- Fixed Styles: Each style required separate computation, limiting real-time applications.
- Global Representation: These methods could not capture local details effectively.

1.1.5 GANs

GAN is a deep learning method in which two neural networks (generator and discriminator) are trained in competition. The generator creates synthetic data while the moderator verifies the truth. This opposing game forces the generator to deliver outputs that closely emulate the dataset. Therefore, GANs are used in image generation, video creation, image and video super-resolution, and generating pictures for all sorts of creative applications; they pose an effective tool for unsupervised learning and data synthesis. The key points of GANs:

- GANs consist of two neural networks, a generator and a discriminator, which are trained to oppose each other.
- The adversarial framework improves the generator's ability to produce data indistinguishable from the real data.
- Applications of GANs include image synthesis, video generation, super-resolution, and even generating audio or text.

Generator: The generator is the neural network that creates new samples of data so that they can appear like real data. This generally starts with random noise as input and produces outputs that become better with constant feedback from the discriminator. The generator aims to create data that can successfully "trick" the discriminator into classifying it as real. Over time, it becomes better at capturing the original data's structure, distribution, and features.

The Journal of Computational Science and Engineering (TJCSE)
ISSN 2583-9055 (Media Online)
Vol 3, No 7, July 2025
PP 112– 129

Discriminator: The discriminator's function is, given the data, to ask whether it is real (from the original training data) or fake (from the generator). Essentially, the critic provides feedback to the generator for its improvements. The model is trained in a way that maximizes its accuracy in separating real data from fake data. Gradually, as the generator improves, the discriminator gets better at spotting more and more minute differences, thereby creating more and more impetus for the generator to refine.

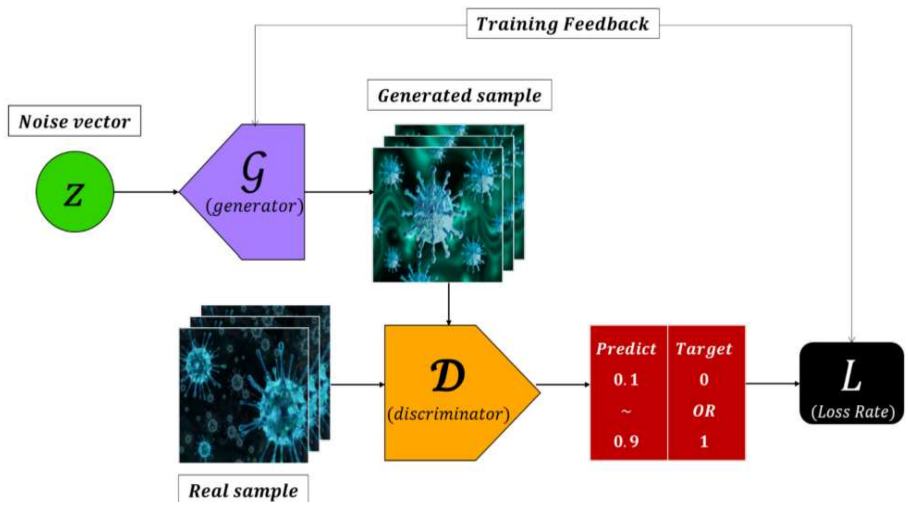

Figure 2: GAN Architecture. (Image Source: [5])

2. Traditional Approaches

2.1 Cycle GAN:

CycleGAN is the first method in unpaired image-to-image translation that solved one of the most significant issues with this domain: dependence on paired datasets. The other approaches to image translation were always based on aligned datasets. CycleGAN, however, can function with unaligned datasets from different domains, such as real-world and cartoon imagery. This rules out the need for paired correspondences, which significantly expands the scope of applications in image-to-image translation. CycleGAN's versatility makes it a perfect application for converting photographs to an artistic style, colorizing, or even translating medical images. It uncouples the reliance on paired data; this makes style transfer incredibly democratized and greatly increases applicability. Since it doesn't rely on the paired datasets, it becomes absolutely perfect where creating datasets aligned with each other becomes impossible or impractical, thus further increasing applicability across diverse domains.

CycleGAN enables image-to-image translation without requiring paired datasets. It uses two generators and two discriminators to transform images from one domain to another, ensuring consistency. For example, Figure 3 can convert photos of horses to zebras and back to horses while maintaining key features. Applications: Widely used in artistic style transfer, medical imaging (e.g., normalizing image datasets), and data augmentation.[6]. CycleGAN does not require strict correspondence between the information of the two spaces, which makes the application of CycleGAN broader. Figure 4 shows the CycleGAN workflow, showing the transfer effect of self-attention and semantic segmentation

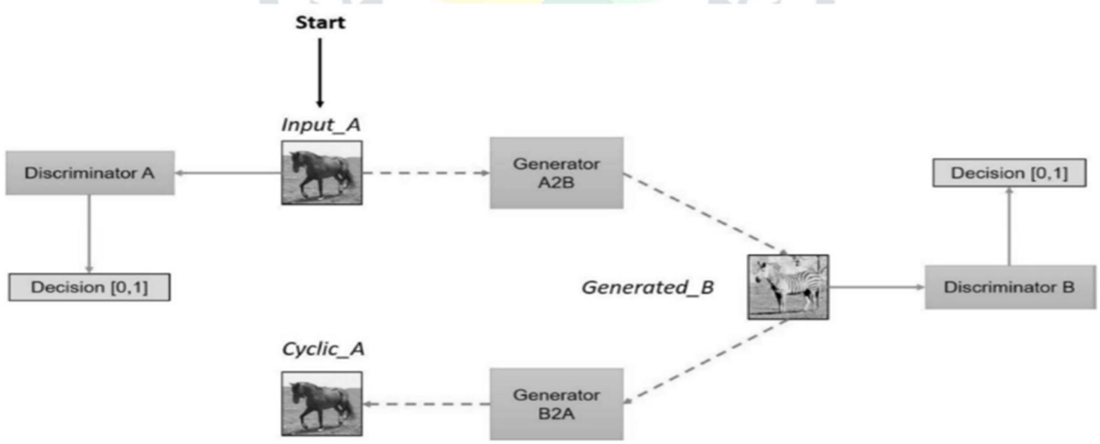

Figure 3: Cycle GAN architecture (Image Source [6])

The algorithm utilizes two fundamental mechanisms to work with: adversarial loss and cycle consistency loss.

- Adversarial Loss: This component ensures that the generated images in the target domain are indistinguishable from real images in that domain. The adversarial training framework consists of two neural networks:
  - ✓ Generator: Creates synthetic images that mimic the target domain.
  - ✓ Discriminator: Evaluates whether an image is real (from the target dataset) or fake (generated by the model).

- Cycle Consistency Loss: This parameter ensures that when an image is transformed into the target domain and back into the original domain, it remains similar to the original input. This helps the model preserve the structural and content integrity of the images during transformation.

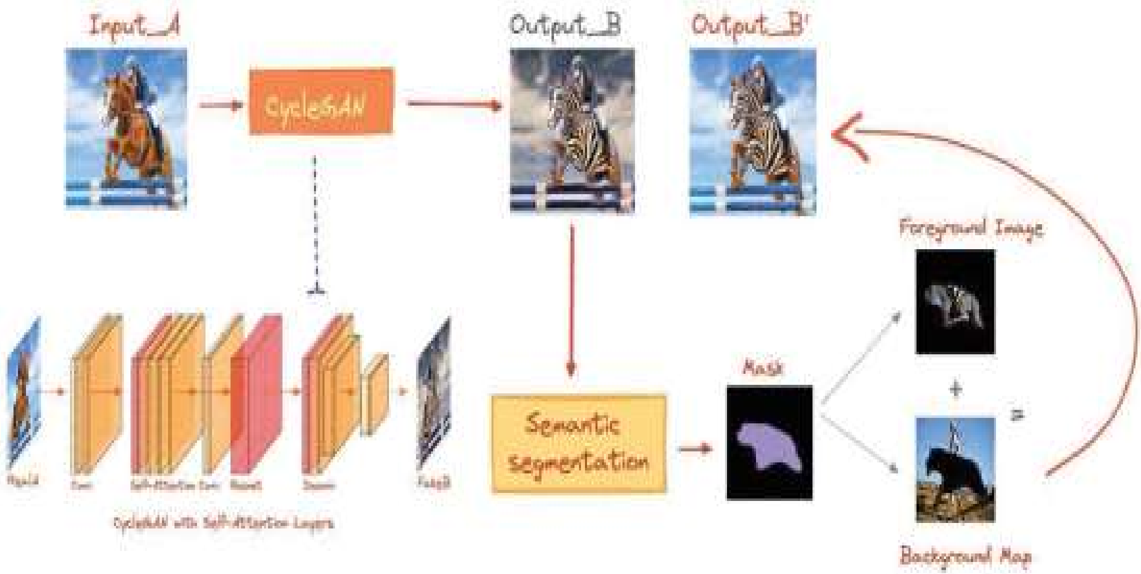

Figure 4. CycleGAN workflow showing the transfer effect of self-attention and semantic segmentation (Image Source [7])

CycleGAN has integrated a variety of algorithms to improve performance and flexibility. Most commonly, ResNet-based generators with skip connections enhance the learning process while conserving some fine details during the translation. Grounded on local image feature patches, PatchGAN discriminators are capable of generating high-quality outputs. Semantic segmentation approaches, like the U-Net architecture, have an encoder-decoder framework that provides skip connections for precise alignment and spatial coherence. The Dynamic Potential Head (DPH) adapts the feature representations, whereas the Dense Instance Affinity Head (DIAH) captures dense pixel-level relationships, thus improving spatial alignment and object-level consistency. Cycle consistency loss, the backbone of CycleGAN, is applied to ensure content preservation by maintaining a bidirectional mapping between image domains. Besides, perceptual loss improves perceptual realism by using features from pre-trained networks. The attention mechanisms and style transfer modules, such as Adaptive Instance Normalization, refine the translation operation by explicitly focusing on salient regions. CycleGAN forms a solid framework for unpaired image style transfer tasks. Datasets such as Cityscapes, Monet2Photo, and Facades are commonly used to train CycleGAN models. These datasets represent diverse domains, enabling applications like converting photographs into artistic paintings, transforming architectural layouts, and changing seasonal landscapes (e.g., summer to winter scenes).

2.1.1 Advancements Over Traditional Methods

- No Need for Paired Data: CycleGAN eliminates the requirement for aligned datasets, making it suitable for domains where paired data is difficult or impossible to obtain.
- Preservation of Content: The cycle consistency loss ensures that the input image's structure and key features are maintained after transformation.
- Flexibility Across Domains: With its ability to handle unpaired datasets, CycleGAN has been applied to diverse areas, including artistic style transfer, domain adaptation, and data augmentation. [8]

Table 1. Comparison of cell detection accuracy (mean ± standard deviation) of CycleGAN, subset-trained CycleGAN (level 2 NIRF), and hierarchical image CycleGAN.

Method	**Precision (%)**	**Recall (%)**	**F1-score (%)**
CycleGAN	82.8 ± 5.7	72.4 ± 8.7	76.9 ± 5.7
CycleGAN	80.0 ± 7.2	77.1 ± 8.1	78.1 ± 4.8
Stratified CycleGAN	80.8 ± 7.6	90.9 ± 6.5	85.0 ± 3.4

2.1.2 Technical challenges of CycleGAN

Despite its ability to perform unpaired image-to-image translation, Cycle GAN faces several technical challenges. One major issue is computational overhead, as training requires two generator-discriminator pairs, effectively doubling the parameters compared to traditional GANs. This increases memory consumption and training time, making it difficult to scale Cycle GAN for high-resolution images. Additionally, the cycle consistency loss, which ensures that an image can be translated back to its original form, adds further computational burden.
To mitigate these challenges, several optimizations have been proposed:

- Efficient Generator Architectures – Using ResNet-based generators with skip connections improves computational efficiency while preserving fine details in translation.

The Journal of Computational Science and Engineering (TJCSE)
ISSN 2583-9055 (Media Online)
Vol 3, No 7, July 2025
PP 112 – 129

- PatchGAN Discriminators – Instead of classifying entire images, PatchGAN focuses on local image patches, significantly reducing computational complexity.
- Hybrid Training Approaches – Techniques such as progressive training allow models to learn from low-resolution images before refining high-resolution outputs, optimizing GPU usage.

Despite these optimizations, CycleGAN still struggles with mode collapse, where the generator produces limited variations of output images. Future research focuses on adaptive loss functions and attention-based mechanisms to enhance robustness while maintaining computational feasibility.

2.2 StyleGAN

StyleGAN is a specialized GAN for creating high-quality, photorealistic images with fine control over style and content. It introduces a style-based generator architecture that allows fine-tuning image attributes, such as facial expressions or artistic styles, by manipulating the latent space. Applications: Used in generating synthetic human faces, artistic image creation, and even data-driven character design. Advances in StyleGAN propel generative modeling to a transformative leap, especially with respect to high-resolution image synthesis. Leveraging dedicated datasets such as fine-tuned facial images in Face-HQ, StyleGAN obtains an impressive combination of photorealism and stylistic diversity, producing high-quality results. The model is preferred for precision and consistency-demanding tasks in detailed and realistic images. [9]

StyleGAN employs a style-based architecture to generate high-fidelity, realistic images with tight control over other aspects like pose, texture, and color. Innovations in StyleGAN2 and StyleGAN3 address the cracks and aliasing with state-of-the-art performance on Frechet Inception Distance (FID) and other metrics. However, StyleGAN is limited by its computational expense and risk of over-training on small datasets; hence, it lacks possibilities for generating diverse outputs without sufficient data. This may include difficulty complying with domain constraints and requires high resource consumption for training and inference.

The Journal of Computational Science and Engineering (TJCSE)
ISSN 2583-9055 (Media Online)
Vol 3, No 7, July 2025
PP 112–129

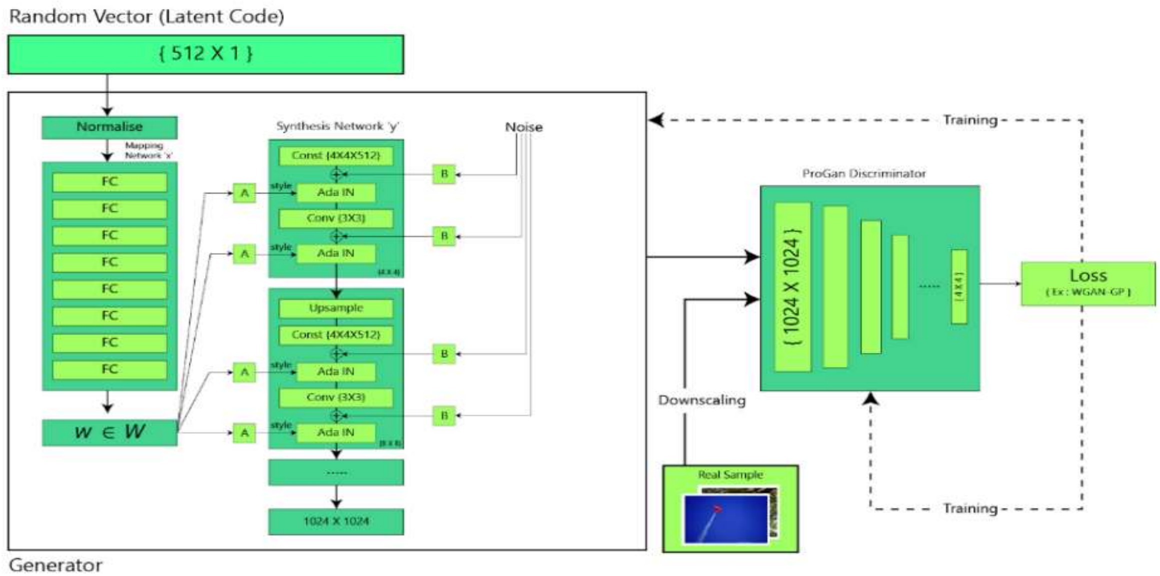

Figure 5: The StyleGAN architecture. (Image Source [2])

2.2.1 StyleGAN's computational challenges:

While StyleGAN is a breakthrough in high-quality image synthesis, it has significant computational challenges. Training requires extensive GPU resources, often taking weeks on high-end hardware, due to its style-based generator architecture and multi-resolution processing. Additionally, StyleGAN's reliance on large-scale datasets poses challenges in domains where high-quality training data is scarce. To address these limitations, researchers have implemented several key improvements:

- StyleGAN2 Optimization – Introduced weight demodulation and path length regularization, reducing artifacts and improving stability during training.
- StyleGAN3 Enhancements – Focused on removing aliasing effects and improving the spatial coherence of generated images, making outputs more realistic.
- Adaptive Data Augmentation – New techniques allow StyleGAN to be trained on smaller datasets by artificially increasing diversity and improving generalization.

Despite these advances, StyleGAN still struggles with real-time applications due to its high computational cost. Future developments aim to reduce inference time and memory usage, enabling broader deployment in applications like content creation, gaming, and medical imaging.

The Journal of Computational Science and Engineering (TJCSE)
ISSN 2583-9055 (Media Online)
Vol 3, No 7, July 2025
PP 112– 129

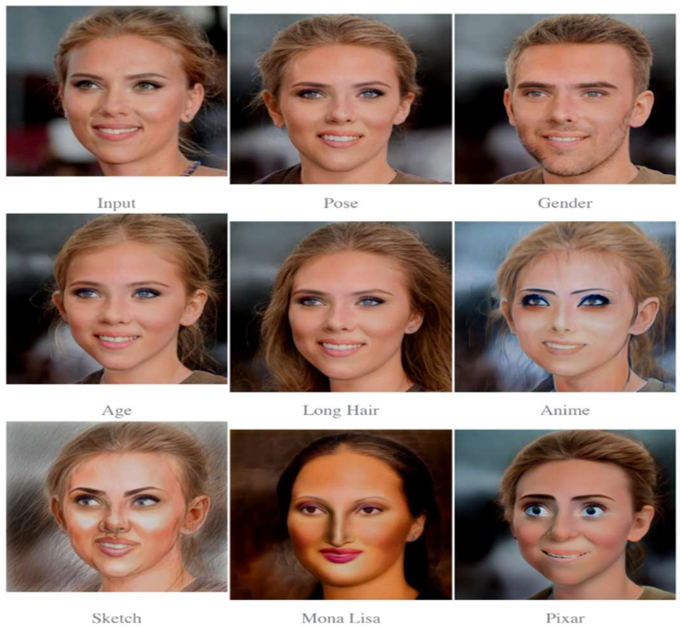

Figure 6: Editing a real image of Scarlett Johansson (on the top left) with StyleGAN. We show both in-domain and out-of-domain manipulations. (Image Source [2])

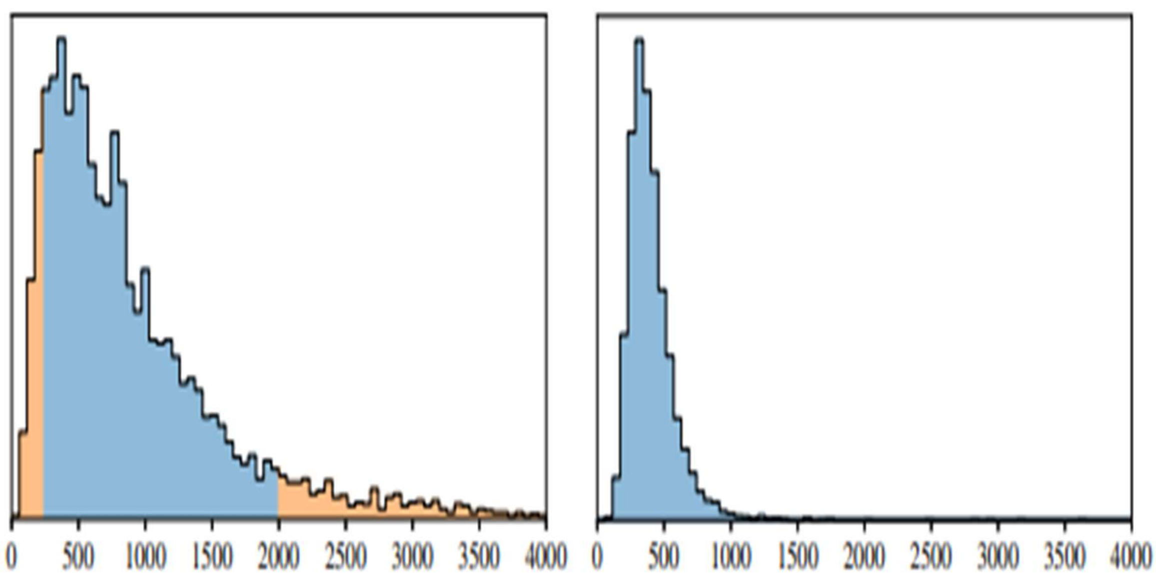

Figure 7: PPL score distribution for an image based on StyleGAN and StyleGAN2

2.3 Neural Style Transfer (NST)

The earliest type of style transfer was NST. The novelty of NST lies in incorporating features of content and style images. This is implemented through pre-trained CNNs, which extract content and style features from corresponding content and style images. While finding a balance between the content-preservation and style-emulation factors, NST produces visual wonders. [8]. Innovations in Realistic Image Synthesis:

The algorithm works by iteratively optimizing an initial noise image to compute both content loss, which ensures structural fidelity, and style loss, which is derived from Gram matrices to replicate the patterns and textures of the target style. Although results are compelling, such an optimization-based approach is computationally costly and slow, and thus not ideal for real-time applications or those demanding fast processing.[10]

Despite its limitations, NST opened the door to later advancements in real-time and adaptive style transfer methods. Its visually pleasing outputs influence artistic and practical applications, including video style transfer, real-time image filters, and creative content generation. NST's legacy lives on as a foundation for more efficient and scalable techniques in style transfer.

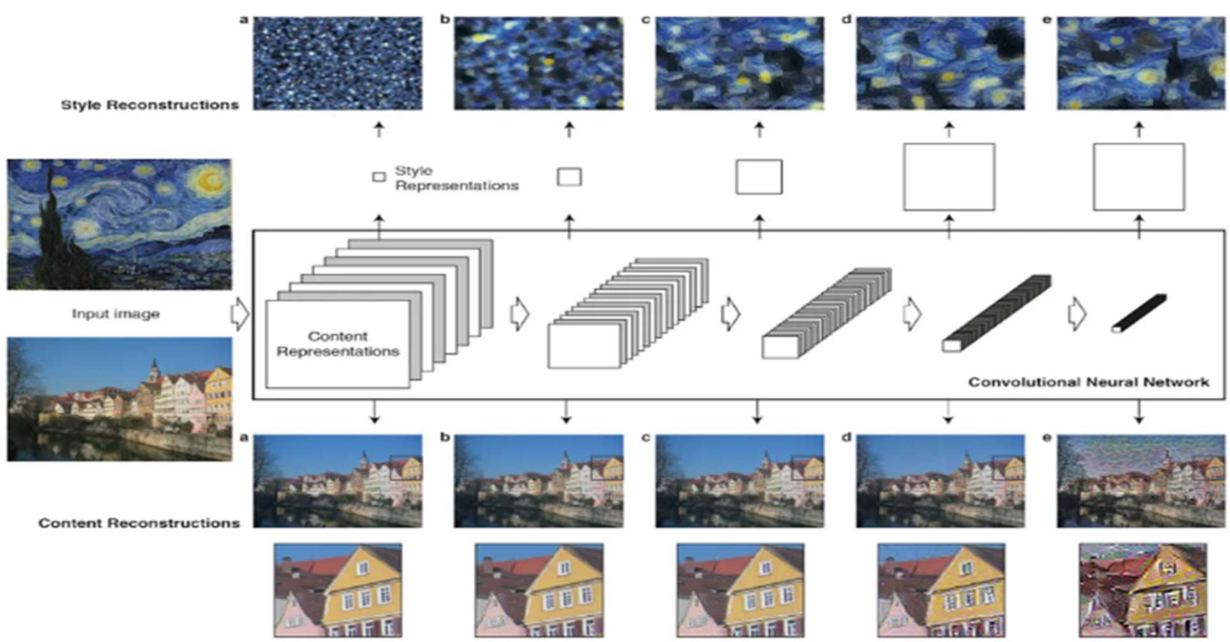

The Journal of Computational Science and Engineering (TJCSE)
ISSN 2583-9055 (Media Online)
Vol 3, No 7, July 2025
PP 112– 129

Figure 8: The NST architecture (Image Source [10])

3. Datasets Used

3.1 MS COCO

The MS COCO (Common Objects in Context)[11] dataset is a large-scale collection of images depicting everyday scenes wherein objects are shown in their natural contexts. It comprises over 80,000 training images annotated for various tasks, including object detection, segmentation, and captioning. This dataset can be a strong resource for such content-based functions because of the varied possible scene-object combinations, providing a broad basis of training material for improving object understanding and contextual relations within the models.

3.2 WikiArt

Wiki Art is an extensive collection of paintings representing different artistic styles, periods, and artists. It contains over 80,000 paintings in categories like Impressionism, Surrealism, and Abstract art. This dataset is particularly used for style-based tasks thanks to its rich variety of textures and patterns that enable explorations of artistic traits and the training of models that understand the nature of styles and how to reproduce them. [12]

3.3 Caltech

The Caltech 101 and Caltech 256 datasets are collections of images organized into 101 and 256 categories of objects, respectively. The Caltech 101 dataset offers approximately 9,000 images, and the Caltech 256 dataset contains around 30,000 images. These datasets are heavily used in style-transfer-based data augmentation tasks and support numerous categories of objects to improve the generalization of the models over different styles and objects.

3.4 Artistic Image Pairs

Artistic Image Pairs datasets consist of paired collections of real and artistic images, such as those used in AnimeGAN and CartoonGAN. These datasets are designed for supervised training in tasks involving actual to cartoon or anime-style transformations, emphasizing keeping the semantic content unchanged while varying stylistic representation; thus, they provide valuable lessons for stylistic transformational tasks while seeking to maintain the underlying meaning behind images.[13]

4. Cycle GAN for Image Style Transfer

CycleGAN is a powerful technique for unpaired image-to-image translation. It allows transformations without aligned datasets and is very effective in real-to-cartoon conversions and day-to-night adaptations. The whole procedure is based on adversarial loss, which guarantees realistic outputs, and cycle consistency loss preserves originality. This allows for stylistic transfer without damaging content to a significant extent.

This permutation of the network has other strengths as well. Cycle consistency enforces the keeping of the essential structure and gives it an added value in artistic renderings and the applications where paired datasets are often a rarity. Nevertheless, CycleGAN implies high computational costs. The dual-control system of the generator-discriminator shifts memory overheads and time, especially on high-resolution images. The existence of the cycle consistency loss adds a load of computation, discreetly pushing it behind real-time style transfer methods.[14] Optimization of this system makes the models much faster, where there are ResNet-based generators with skip connections for detail retention, and the Patch-GAN discriminators [15] decrease the computation burden. However, articulation and global coherence problems are challenging to solve, and some networks cause artifacts. Architectures based on semantic segmentation (such as U-Net) could fit much [16] better for aligning this structure, although they haven't made their way into the realm of CycleGAN. The DIAH and DPH techniques improve feature adaptation but also increase computational cost. A significant trade-off is between realism and fidelity with CycleGAN [17]. Cycle consistency prevents losing content, but the fine stylistic details are uncatchable. Future improvements should optimize the losses used, integrate attention into the process, and use lightweight architectures to keep efficiency up while delivering high-quality results.

4.1 Quantitative Comparison of Style Transfer Methods

Different style transfer models vary in [18] speed, computational efficiency, and scalability:

Method	Execution Time	GPU/Memory Usage	Scalability

CycleGAN	**Moderate** (Slower than real-time methods)	**High** (Two-generator-discriminator pairs)	**High** (No paired dataset required)
StyleGAN	**Very Slow** (Days to train, seconds to generate)	**Very High** (Multiple GPUs needed)	**Moderate** (Best for large datasets)

5. Conclusion

CycleGAN has revolutionized areas of image style transfer by overcoming limitations of traditional style transfer methods that depend on paired datasets or iterative optimization. Indeed, these conventional approaches lack flexibility, computational efficiency, and applicability in real-life scenarios. CycleGAN avoids these through an ingenious approach rooted in unpaired datasets and cycle consistency loss to guarantee that the content-coding phase is preserved through the transformation. This massive boost toward style transfer is now conceivable even when no synchronized training data is available or feasible.

Employing adversarial loss, CycleGAN ensures a high visual quality of the generated images, while cycle consistency loss guarantees that the transformation does not distort the original structure. Such granularity makes unprecedented domain adaptation over such tasks as photo-to-painting conversion and real-world photo-to-cartoon transformation or seasonal scenery alteration possible, endearing CycleGAN to a broader group because of the associated diversity and versatility- the applications cut across art, health, and wellness, and machine learning for data augmentation.

CycleGAN, with its advantages, faces some computational hurdles. Thus, any overlap in content and style features on the domains might bring some artifacts to the images composed, which might question the outcome's realism. Further, highly demanding resource models take time for training intervals, making real-time usages impossible in resource-constrained settings. Work is to consider reducing computational overhead, combining attention techniques for better spatial consistency, and optimizing training architectures in future work to enhance efficiency.

CycleGAN has revolutionized style transfer, finding plenty of room for generalization and versatility, which remains unrivaled to this day. The range of its dependability in pursuing many domains, along with retaining content structure, makes CycleGAN an absolute favorite among researchers and practitioners. CycleGAN's future lies in being a core bedrock in driving innovations around AI-based artistic rendering and content automation, among many, in medical imaging applications. CycleGAN solves some cardinal computational obstacles and paves the way for the following scalable steps in joining artistic creativity to harmonious technological accomplishments. CycleGAN can further bridge the gap between artistic creativity and technological advancements, shaping the future of visual computing.

References

1. Gatys, L. A., Ecker, A. S., & Bethge, M. (2015). Image Style Transfer Using Convolutional Neural Networks. Computer Vision and Pattern Recognition (CVPR), 2016,
2. R. G. A. M. N. Amit H. Bermano, “Methods and Applications of StyleGAN,” Methods and Applications of StyleGAN, 2024.
3. Chen, J., Yang, G., Zhao, H., & Ramasamy, M. (2020). Audio style transfer using shallow convolutional networks and random filters. Multimedia Tools and Applications, 79(21), 15043-15057.
4. Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2242-2251.
5. Park, S. W., Ko, J. S., Huh, J. H., & Kim, J. C. (2021). Review on generative adversarial networks: focusing on computer vision and its applications. Electronics, 10(10), 1216.
6. Johnson, J., Alahi, A., & Fei-Fei, L. (2016). Perceptual Losses for Real-Time Style Transfer and Super-Resolution. European Conference on Computer Vision (ECCV).
7. Yan, Hang. “A New CycleGAN-Based Style Transfer Method.” A New CycleGAN-Based Style Transfer Method, 2023, p.8
8. Wang, H., & Davis, P. L. (2017). Advancements in Bidirectional Image Transformation using Cycle GANs. IEEE Transactions on Image Processing, 26(8), 3902-3914.(Reference: COS598B_spr2018_NeuralS…)

9. Karras, T., Laine, S., & Aila, T. (2019). A Style-Based Generator Architecture for Generative Adversarial Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, 42(12), 2842-2857.
10. Gatys, L. A., Ecker, A. S., & Bethge, M. (2016). Image style transfer using convolutional neural networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 2414-2423) Zhang, Y., & Robinson, S. (2018). Bidirectional Image-to-Image Translation: Applications and Challenges with Cycle GANs. ACM Transactions on Graphics, 35(4), 127-1:127-15. (Reference: Image Neural Transfer)
11. García-Domínguez, M., Domínguez, C., Heras, J., Mata, E., & Pascual, V. (2021). Neural Style Transfer and Unpaired Image-to-Image Translation to Deal with the Domain Shift Problem on Spheroid Segmentation.
12. Carter, G., & Smith, T. (2022). Enhancing Artistic Expression with CycleGAN: Applications and Techniques.
13. Praramadhan, A. A., & Saputra, G. E. (2021). Cycle Generative Adversarial Networks Algorithm with Style Transfer for Image Generation.
14. Li, S., & Miller, J. (2021). Bidirectional Image-to-Image Translation for Creative Applications: A Cycle GAN Approach. Journal of Computer Science and Technology,
15. Alotaibi, Aziz. (2020). Deep Generative Adversarial Networks for Image-to-Image Translation: A Review. Symmetry. 12. 1705. 10.3390/sym12101705.
16. Kim, E., & Foster, L. (2017). Transformative Potentials of Cycle GANs in Creative Image Synthesis. Computers & Graphics, 94-107.
17. Jin, X., Qi, Y., & Wu, S. (2017). CycleGAN Face-off.
18. Karras, T. (2019). A Style-Based Generator Architecture for Generative Adversarial Networks. arXiv preprint arXiv:1812.04948.