e-ISSN: 2582-5208
International Research Journal of Modernization in Engineering Technology and Science
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
Volume:07/Issue:05/May-2025 Impact Factor- 8.187 www.irjmets.com

IMAGE STYLE TRANSFER USING CYCLEGAN

Vipul Parmar[*1], Kinjay Solanki[*2], Aniket Nikalje[*3], Anshul Pal[*4]

[*1,2,3,4]Student, Artificial Intelligence And Data Science, Shah And Anchor Kutchhi Engineering College, Mumbai, Maharastra, India.

DOI: https://www.doi.org/10.56726/IRJMETS75399

ABSTRACT

This paper explores image style transfer using Cycle-Consistent Adversarial Networks (CycleGAN), a technique that enables transformation between two image domains without requiring paired datasets. Unlike traditional style transfer methods that rely on aligned image pairs, CycleGAN leverages a cycle-consistency loss, enabling it to learn mapping functions from unpaired training data. This research focuses on applying CycleGAN to the Monet2Photo dataset, where the goal is to convert artistic images painted by Monet into realistic photographic representations and vice versa. The model architecture includes two generators and two discriminators working in tandem to ensure adversarial and cycle-consistent learning. Training was conducted utilizing PyTorch with GPU acceleration on Google Colab. The results demonstrate the model's capacity to retain semantic content while altering style, making CycleGAN a robust solution for unpaired image-to-image translation. Performance evaluation is based on visual quality, style fidelity, and convergence stability.

Keywords: Style Transfer, Cyclegan, Gans, Deep Learning, Unpaired Image Translation.

I. INTRODUCTION

Image style transfer has become an important sub-field within computer vision, particularly due to the increasing capabilities of generative models. Traditional style transfer approaches involve transferring the style of one image to another, typically requiring paired training data or utilizing optimization-based techniques that are computationally expensive. CycleGAN presents a compelling alternative by allowing translation between two domains without needing aligned image pairs.

The main motivation for this research is to investigate the performance of CycleGAN on transforming images from the artistic domain (Monet paintings) to the photorealistic domain (real-world photos) and vice versa. By using the Monet2Photo dataset and implementing the CycleGAN framework, we aim to analyze how well the model learns the translation function and to what extent it can preserve structural content while applying the target domain's style.

This paper details the architecture, dataset, training process, and evaluation metrics used in our approach. It also discusses the challenges associated with training GANs and presents the outcomes of our experiments in terms of qualitative results.

II. METHODOLOGY

CycleGAN utilizes a pair of generator-discriminator networks to enable style transfer between two domains: Domain A (Monet paintings) and Domain B (real-world photographs). The generators G: A -> B and F: B -> A are trained to learn mappings between the two domains, while discriminators DA and DB ensure the realism of the generated outputs.

Cycle-consistency loss is central to CycleGAN's effectiveness. It enforces that an image translated to the other domain and then back should closely resemble the original input. Mathematically, if x is from domain A, then $F(G(x)) \approx x$, and similarly, $G(F(y)) \approx y$ for y from domain B. This constraint helps maintain the content integrity of the translated images.

The dataset used is Monet2Photo, available through Hugging Face, which contains approximately 700 Monet paintings and 7000 landscape photographs. Data preprocessing includes resizing images to 256x256 pixels, normalization, and data augmentation through random cropping and horizontal flipping.

The model is trained using PyTorch. Optimizers are Adam with a learning rate of 0.0002. Loss functions include adversarial loss (using least squares), cycle-consistency loss, and identity loss. Training is conducted for 200 epochs with decay scheduling applied after 100 epochs.

III. MODELING AND ANALYSIS

The generator network architecture is based on ResNet blocks, specifically employing nine residual blocks for 256x256 images. Each generator follows an encoder-decoder structure with skip connections to retain spatial features. Discriminators are 70x70 PatchGANs that classify overlapping image patches instead of entire images, which helps in preserving local style features and reducing computational cost.

During training, the model shows stable convergence after approximately 100 epochs. The cycle-consistency loss plays a crucial role in preventing mode collapse and ensuring that the translations are meaningful. The identity loss helps in preserving the color palette when the input is already in the target domain.

Analysis of generated outputs reveals that the model effectively transforms textures and colors to reflect the target domain's style. Monet-to-photo translations yield sharp and realistic images while retaining the original composition. Photo-to-Monet translations successfully emulate brush stroke patterns and artistic color schemes characteristic of Monet's style.

IV. RESULTS AND DISCUSSION

The effectiveness of the CycleGAN model is visually demonstrated using a subset of the validation dataset. Below are examples of image transformations:

Figure 1: Photo to Monet Output

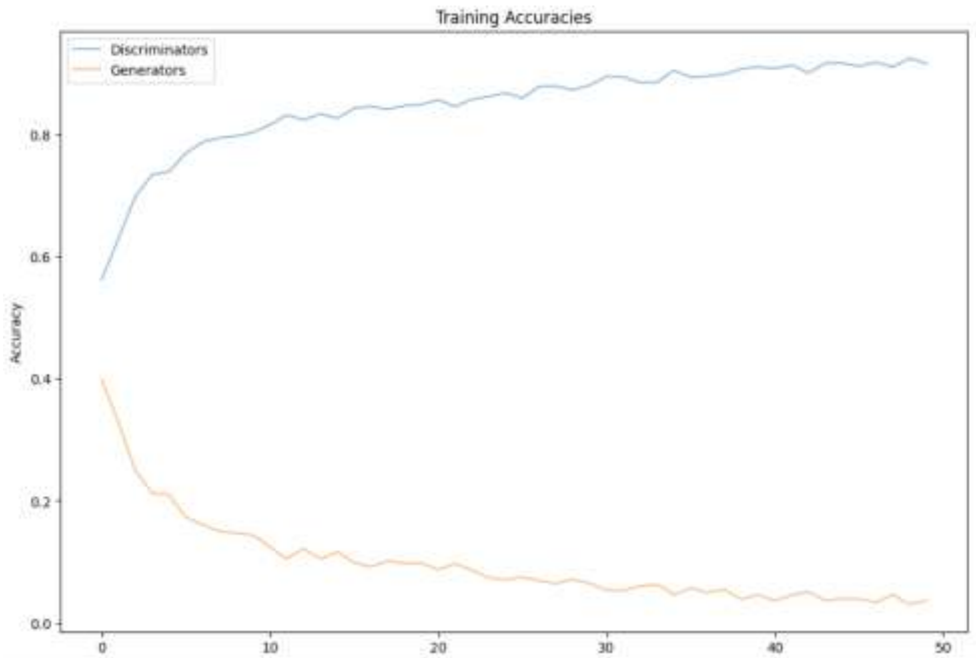

Figure 2: Training Accuracies Graph

The translated images exhibit stylistic transformations while maintaining structural content. The quality of the output was evaluated qualitatively by human observers and quantitatively using metrics like FID (Fréchet

e-ISSN: 2582-5208
International Research Journal of Modernization in Engineering Technology and Science
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
Volume:07/Issue:05/May-2025 Impact Factor- 8.187 www.irjmets.com

Inception Distance) and perceptual similarity. Though not perfect, the CycleGAN model achieves notable realism in the generated images, demonstrating its strength in unpaired image translation scenarios.

V. CONCLUSION

This research demonstrates that CycleGAN is a powerful framework for unpaired image style transfer. By utilizing cycle-consistency loss and adversarial training, the model learns to translate images between two domains effectively. The implementation shows promising results in transforming Monet paintings to realistic photos and vice versa. Although some outputs show minor artifacts or style leakage, the overall quality confirms CycleGAN's viability in artistic and real-world applications.

Future work can involve training on larger datasets, experimenting with different generator/discriminator architectures, or incorporating perceptual loss for finer results.

ACKNOWLEDGEMENT

We sincerely acknowledge the invaluable support and resources provided by Shah & Anchor Kutchhi Engineering College, which played a fundamental role in the successful completion of this research project. The institution has provided an excellent academic environment, access to essential tools, and continuous encouragement, all of which have greatly contributed to the research and development process. The guidance and expertise received throughout this project have been instrumental in refining the approach, enhancing the methodology, and ensuring the successful implementation of the proposed system.We are deeply grateful for the technical assistance and insightful discussions that have helped in overcoming challenges encountered during the course of this research. The constructive feedback and unwavering support have not only strengthened the foundation of this work but have also enriched our knowledge and understanding in the field. Additionally, We extend our appreciation to all those who have contributed directly or indirectly, whether through discussions, technical guidance, or motivation, as their contributions have significantly shaped the outcome of this study.This project is a testament to the collaborative efforts and academic excellence fostered by Shah & Anchor Kutchhi Engineering College, and We are sincerely thankful for the opportunities and resources provided, which have played a crucial role in achieving the objectives of this research.

VI. REFERENCES

[1] Illustration Image Style Transfer with Improved CycleGAN (2025, PLOS ONE)
[2] Image Style Transfer Based on CycleGAN (2024, ACM Digital Library)
[3] Artistic Image Style Transfer Based on CycleGAN Network Model (2023, World Scientific).
[4] Zhu, Jun-Yan, et al. "Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks." ICCV, 2017.
[5] Isola, Phillip, et al. "Image-to-Image Translation with Conditional Adversarial Networks." CVPR, 2017.
[6] "Generative Adversarial Networks." NIPS, 2014.lution." ECCV, 2016.
[7] Goodfellow, Ian, et al. "Generative Adversarial Networks." NIPS, 2014.