REAL-WORLD PICTURES TO ARTISTIC PAINTINGS STYLE TRANSFER USING GAN

SHREYAS PUTTARAJU

Introduction to Style Transfer

Artistic style transfer using deep learning techniques, particularly the CycleGAN architecture, enables individuals to transform their photographs into paintings, fostering creative expression and democratizing art.

OBJECTIVES

01

The project utilizes GANs, specifically the CycleGAN architecture, to convert photographs into artistic paintings.

02

The method involves two generators and two discriminators in a cycle-consistent framework to achieve realistic and artistically consistent outputs.

03

Training the model requires a dataset of unpaired images, consisting of photographs and paintings.

04

Experimental results demonstrate the successful reproduction of artistic styles on photographs, enabling users to create unique artworks and experiment with different styles.

DATA

Data consists of 50 Monet Paintings and 50 Real World Photos randomly downloaded from Google.

All the photos were colored images i.e., had 3 color channels

Images were pre-processed - Scaled and normalized before training

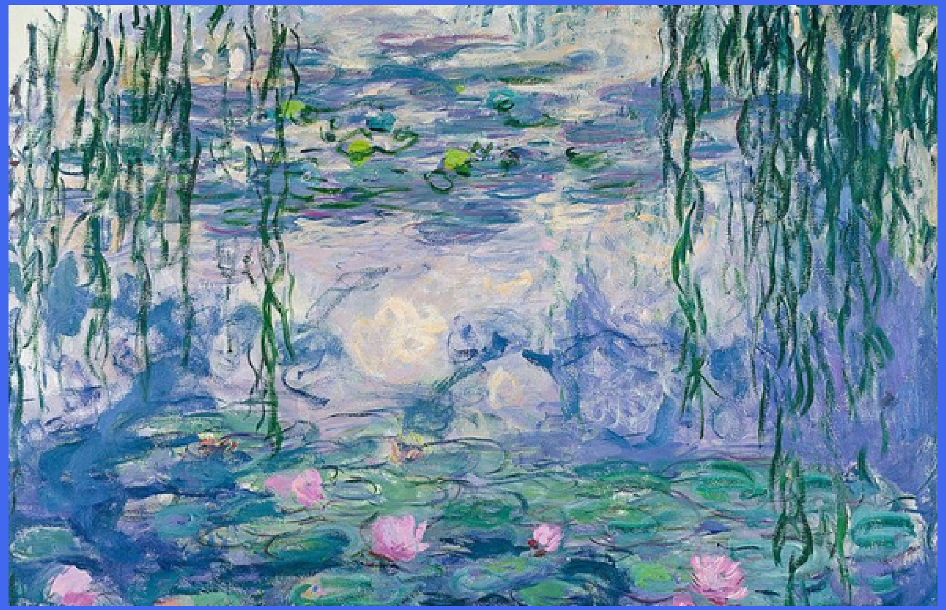

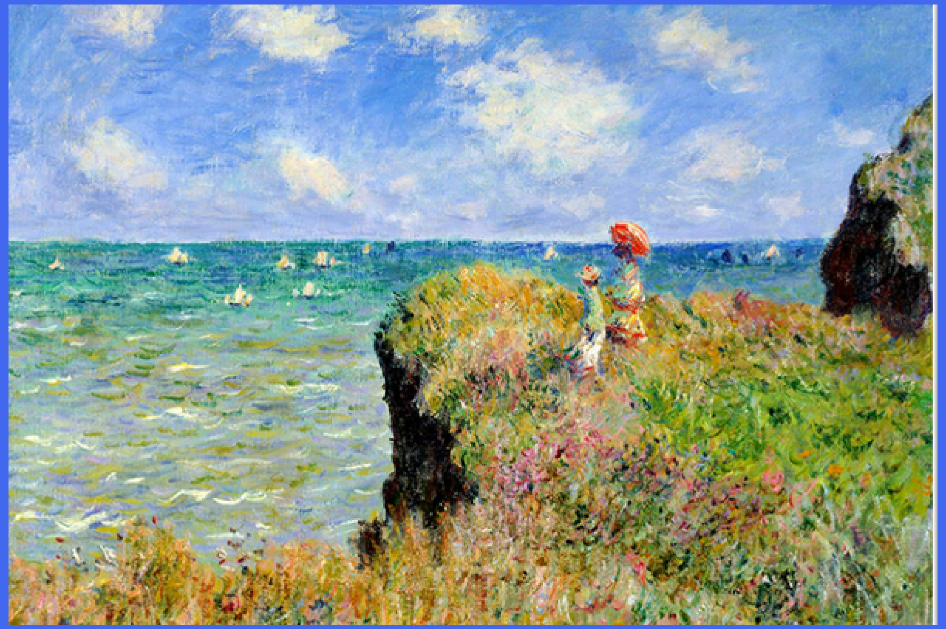

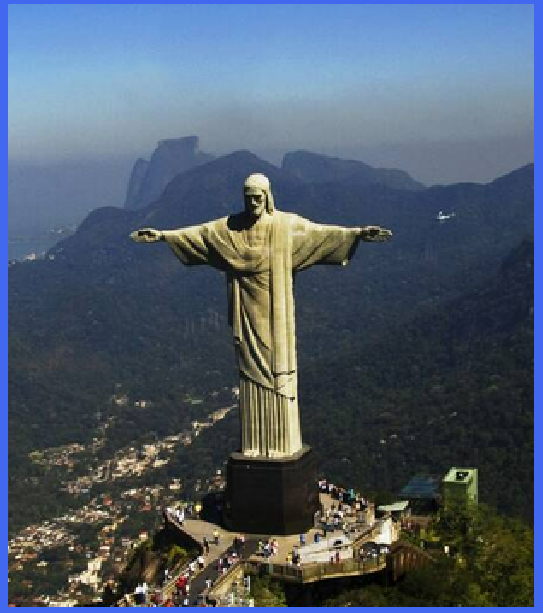

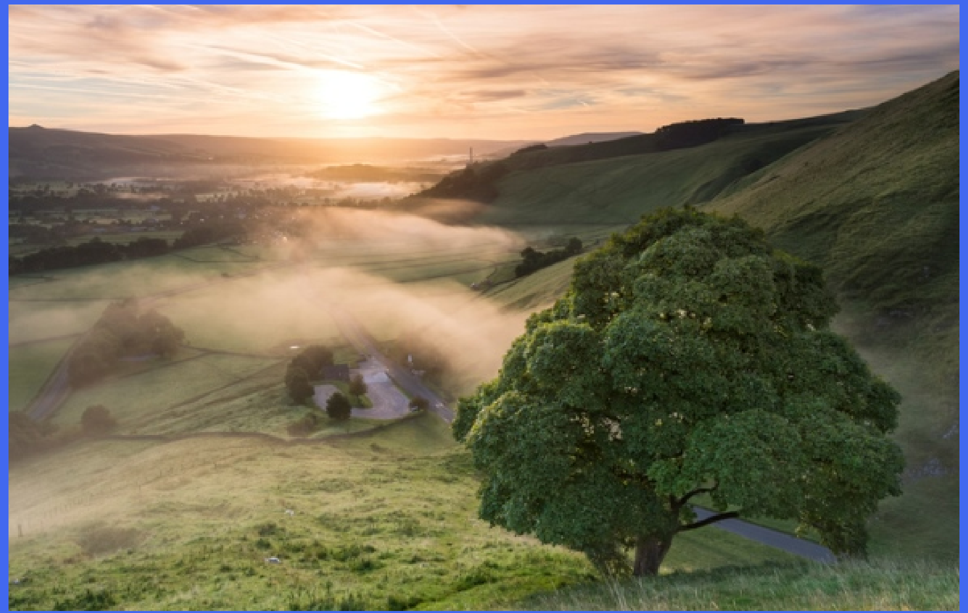

Methodology

01 MODEL BUILDING

02 TRAINING

03 LOSS FUNCTIONS

04 EVALUATE RESULTS

MODEL BUILDING

01

GAN Background

A Generative Adversarial Network (GAN) is a deep learning model that trains two neural networks, a generator and a discriminator, to generate new data samples that are indistinguishable from real data.

The generator network takes in a random noise input and maps it to a synthetic data output through a series of operations in the network.

A discriminator network is trained to distinguish between real and generated data and used in combination with a generator to improve the quality of generated data.

Generally, GAN's work on labeled data.

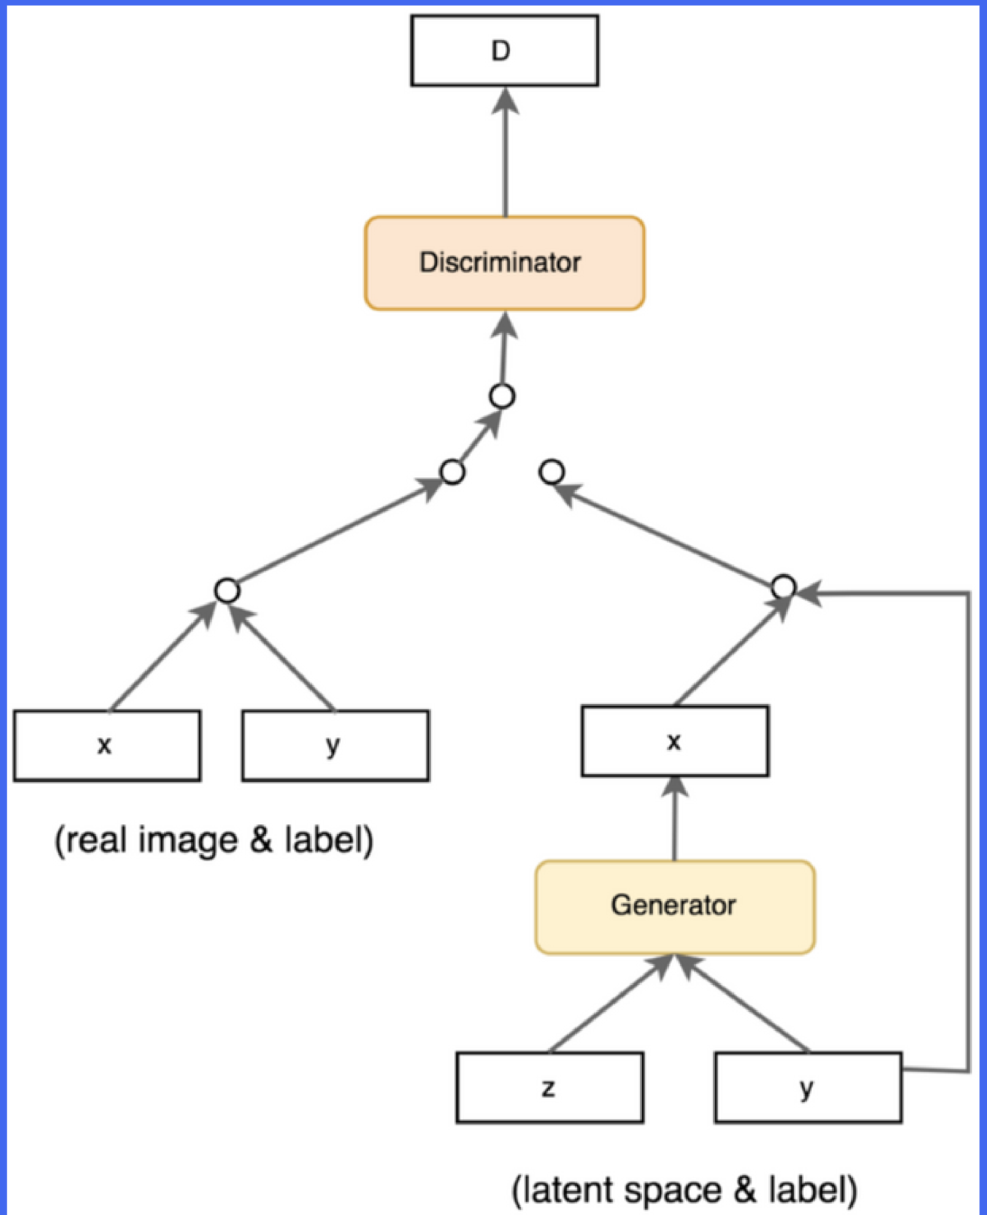

MODEL BUILDING

01

CYCLE-GAN

Cycle-GAN is a GAN architecture for image-to-image translation tasks where the goal is to learn a mapping between two image domains.

It uses two generators, one for mapping source images to target images and another for mapping target images back to source images, and two discriminators to ensure that the generated images are realistic.

Cycle-GAN can be trained without the need for paired training data.

We train the Cycle-GAN to convert real photos to Monet Paintings.

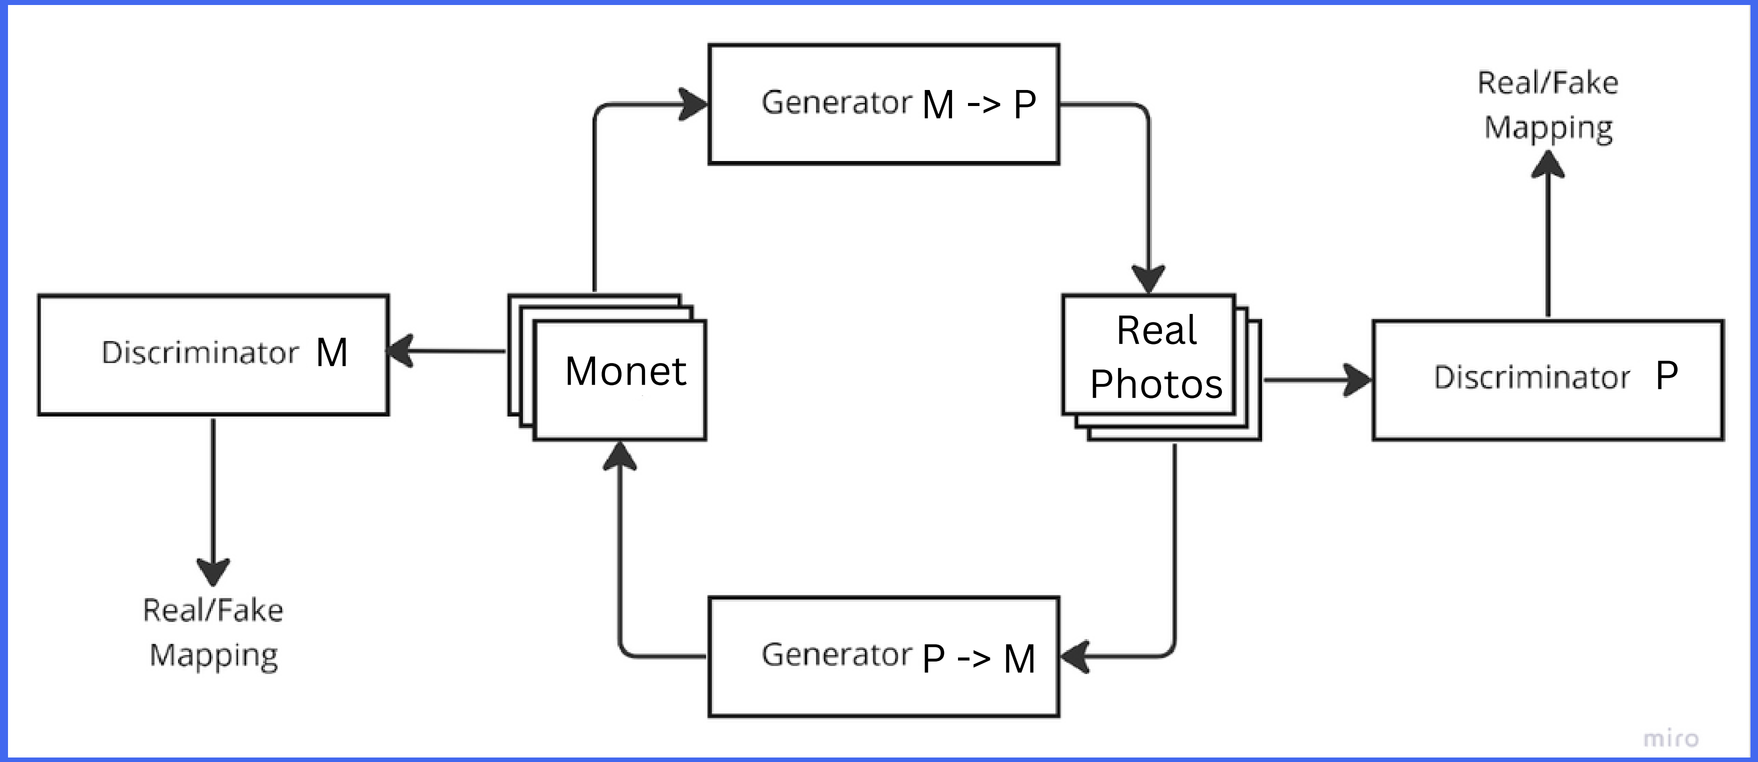

TRAINING

02

After the model was built, the model was trained for 250 epochs with batch size of 1 image.

This was suggested in the original paper, but Google Collab was terminating after 5-6 hours.

We were not able to get the expected results as the model would have not learnt much in 50-60 epochs.

So we show the results from the original paper in the presentation. Our results are in the final report that was submitted last week.

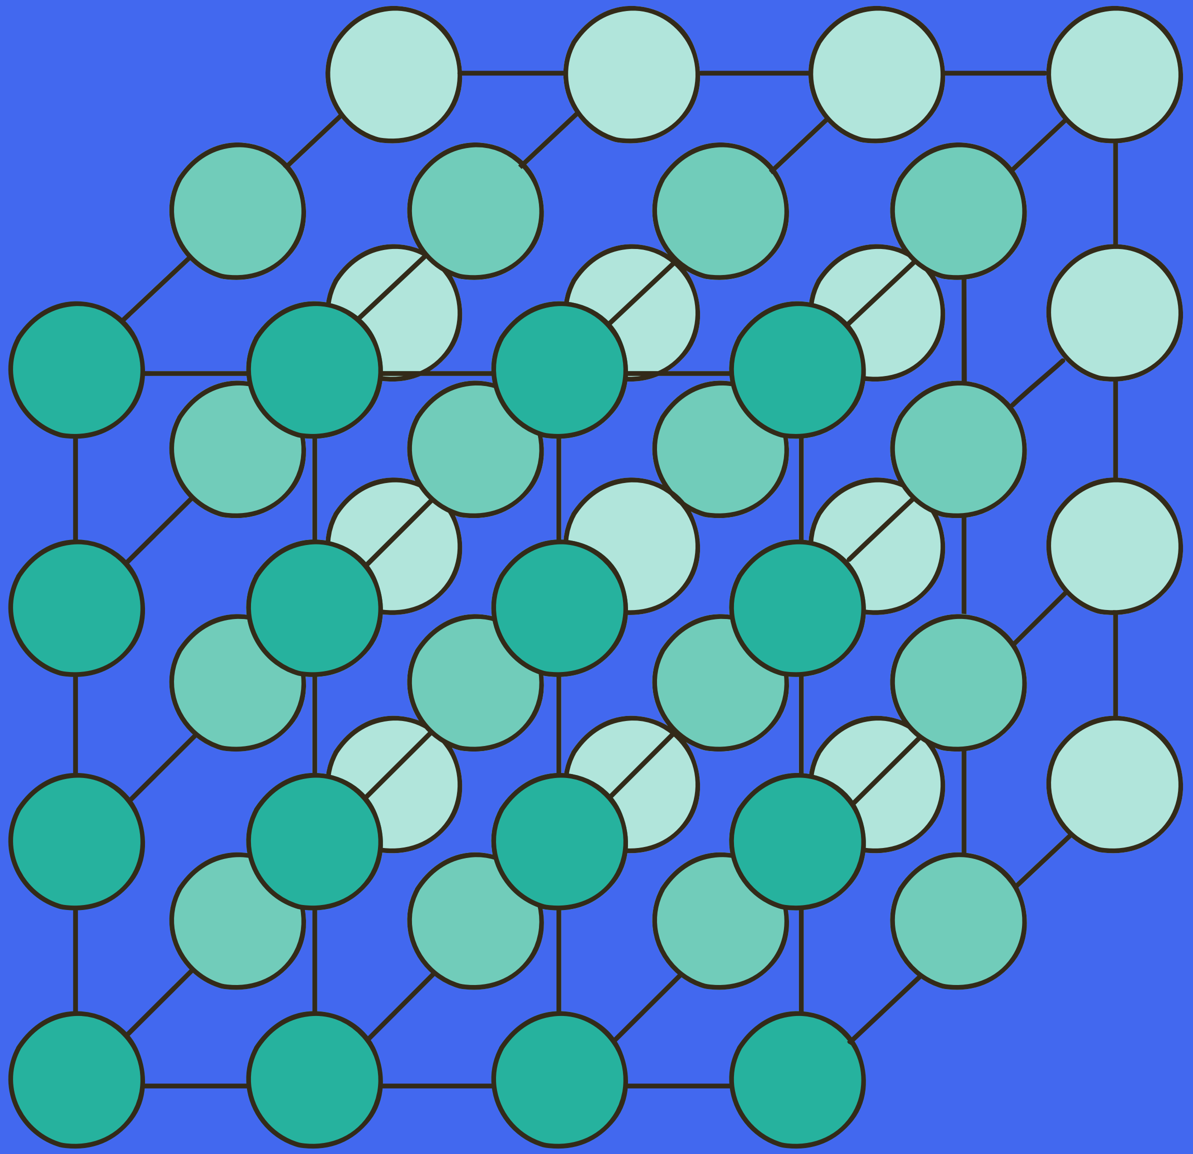

LOSS FUNCTIONS

Aadversarial loss

The adversarial loss encourages the generator network to produce images that are similar to the target domain, while the discriminator network tries to distinguish between the generated and real target images.

This is achieved through a minimax game between the two networks, where the generator tries to minimize the loss while the discriminator tries to maximize it.

The adversarial loss is typically calculated as the binary cross-entropy loss between the predictions of the discriminator network and the target labels.

The target labels are either "real" or "fake", depending on whether the input to the discriminator is a real target image or a generated image.

LOSS FUNCTIONS

03

Cycle Consistency Loss

The cycle consistency loss function is used to enforce the property that a source image should remain unchanged after being translated to the target domain and back to the source domain.

It ensures that the generator network is able to maintain the content of the source image after translation.

The cycle consistency loss is calculated as the mean absolute difference between the original source image and the reconstructed source image obtained by translating the generated target image back to the source domain using a second generator network.

The second generator network is trained to translate target images back to source images.

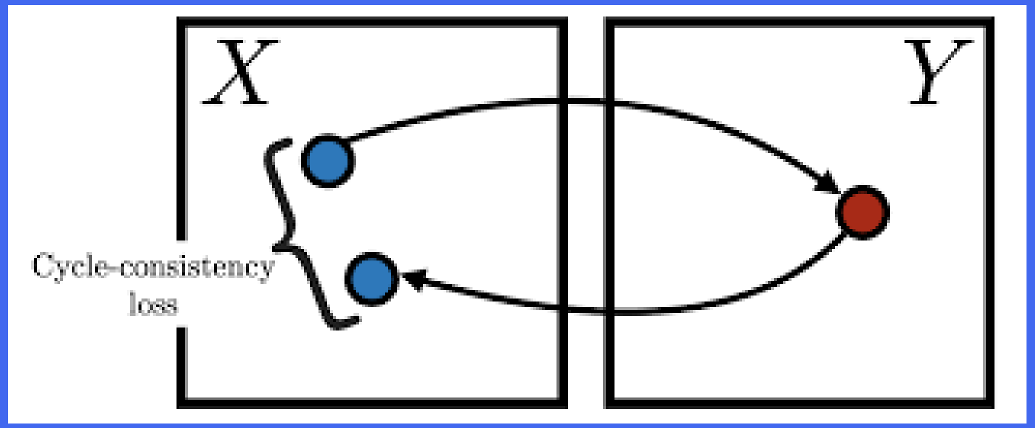

EVALUATE RESULTS

Democratization of art: Enables non-artists to create personalized artworks and explore various styles.

Generalizability of CycleGAN: Successfully transfers artistic styles to different real-world photographs.

Artistic consistency: Captures style traits of input paintings, preserving the integrity of the original style.

Advancing style transfer research: Improves CycleGAN architecture and offers insights for further development and expansion.

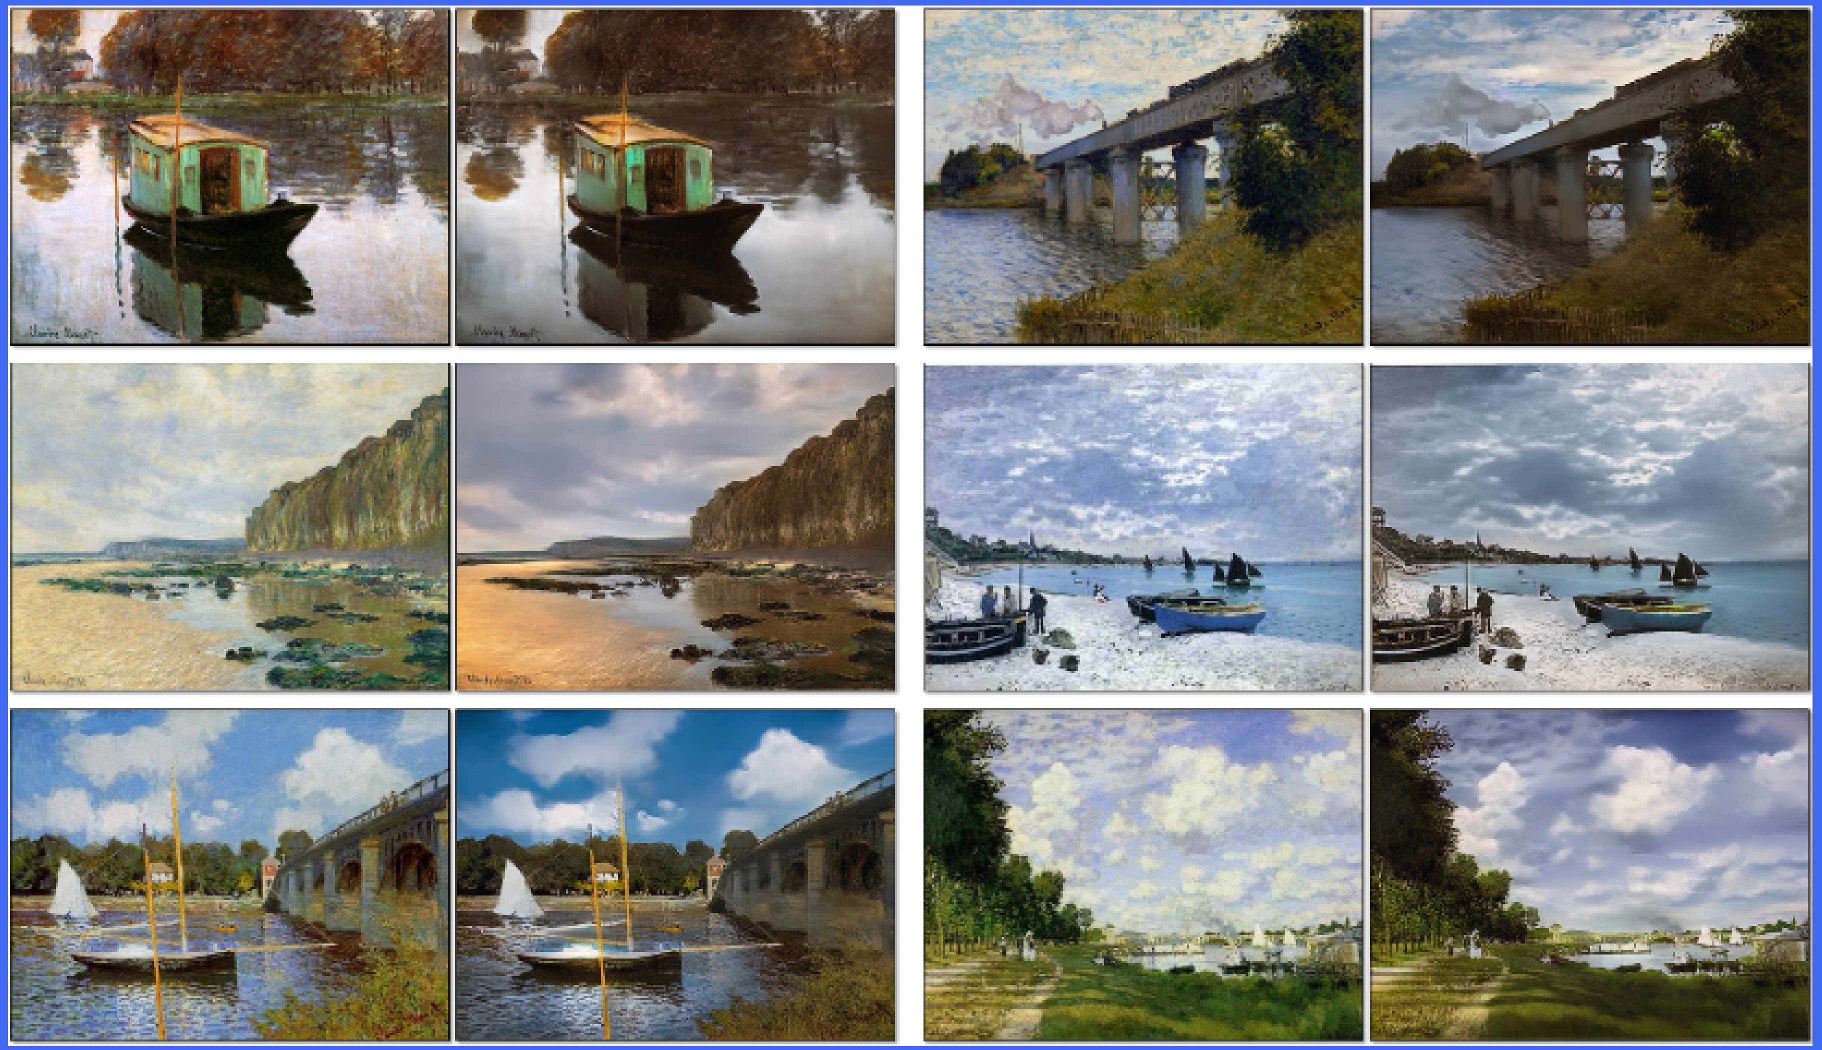

CONCLUSION

The study demonstrates the effectiveness of the CycleGAN architecture in transforming realistic images into artistic paintings, retaining the essence of the original photographs while showcasing visually appealing styles. These findings highlight the significance of democratizing art and promoting research in the intersection of technology and creative expression.

REFERENCES

1) Generative Adversarial Networks Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David WardeFarley, Sherjil Ozair, Aaron Courville, Yoshua Bengio https://arxiv.org/abs/1406.2661

2) Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks Jun-Yan Zhu, Taesung Park, Phillip Isola, Alexei A. Efros; Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2223-2232 https://openaccess.thecvf.com/content_iccv_2017/html/Zhu_Unpaired_ImageToImage_Translation_ICCV_2017_paper.html

THANK YOU!