

## Revitalizing Art with Technology: A Deep Learning Approach to Virtual Restoration

Nurrohmah Endah Putranti <sup>(1)</sup>, Shyang-Jye Chang <sup>(2)</sup>, Muhammad Raffiudin <sup>(3)\*</sup>

<sup>1,2</sup>Department of Mechanical Engineering, College of Engineering, National Yunlin University of Science & Technology, Yunlin, Taiwan

<sup>3</sup>Computer Science Department, Faculty of Science, Chiang Mai University, Chiang Mai, Thailand

e-mail : {m11111073, changjye}@yuntech.edu.tw, muhammad\_r@cmu.ac.th.

\* Corresponding author.

This article was submitted on 20 October 2024, revised on 7 November 2024, accepted on 7 November 2024, and published on 31 January 2025.

### Abstract

*This study evaluates CycleGAN's performance in virtual painting restoration, focusing on color restoration and detail reproduction. We compiled datasets categorized by art styles and conditions to achieve accurate restorations without altering original reference materials. Various paintings and those with a yellow filter, to create effective training datasets for CycleGAN. The model utilized cycle consistency loss and advanced data augmentation techniques. We assessed the results using PSNR, SSIM, and Color Inspector metrics, focusing on Claude Monet's *Nasturtiums in a Blue Vase* and Hermann Corrodi's *Prayers at Dawn*. The findings demonstrate superior color recovery and preservation of intricate details compared to other methods, confirmed through quantitative and qualitative evaluations. Key contributions include employing CycleGAN for art restoration, model evaluation, and framework development. Practical implications extend to art conservation, digital library enhancement, art education, and broader access to restored works. Future research may explore dataset expansion, complex architectures, interdisciplinary collaboration, automated evaluation tools, and improved technologies for real-time restoration applications. In conclusion, CycleGAN holds promise for digital art conservation, with ongoing efforts aimed at integrating across fields for effective cultural preservation.*

**Keywords:** Art Restoration, CycleGAN, Deep Learning, PSNR, SSIM

### Abstrak

Penelitian ini mengevaluasi kinerja CycleGAN dalam restorasi lukisan virtual, dengan fokus pada pemulihan warna dan detail. Kami menyusun dataset yang dikategorikan berdasarkan gaya seni dan kondisi untuk mencapai restorasi akurat tanpa mengubah materi referensi asli. Beberapa lukisan, termasuk yang dengan filter kuning, dirusak untuk membuat dataset pelatihan yang efektif bagi CycleGAN. Model ini memanfaatkan cycle consistency loss dan teknik augmentasi data canggih. Hasil dievaluasi menggunakan metrik PSNR, SSIM, dan Color Inspector, dengan fokus pada *Nasturtiums in a Blue Vase* karya Claude Monet dan *Prayers at Dawn* karya Hermann Corrodi. Temuan menunjukkan pemulihan warna yang unggul dan pelestarian detail halus dibandingkan metode lain, yang dikonfirmasi melalui evaluasi kuantitatif dan kualitatif. Kontribusi utama termasuk penggunaan CycleGAN untuk restorasi seni, evaluasi model, dan pengembangan kerangka kerja. Implikasi praktis mencakup konservasi seni, peningkatan perpustakaan digital, pendidikan seni, dan akses yang lebih luas terhadap karya yang direstorasi. Penelitian di masa depan dapat mengeksplorasi perluasan dataset, arsitektur kompleks, kolaborasi lintas disiplin, alat evaluasi otomatis, dan teknologi untuk aplikasi restorasi waktu nyata. Kesimpulannya, CycleGAN menunjukkan potensi dalam konservasi seni digital, dengan upaya terus berlanjut untuk integrasi dengan bidang lain guna pelestarian budaya yang efektif.

**Kata Kunci:** Restorasi Seni, CycleGAN, Deep Learning, PSNR, SSIM



## 1. INTRODUCTION

Preserving the intrinsic value of artwork is a primary goal of art restoration. Art collectors often assign lower value to paintings with damage, such as discoloration, dullness, yellowing, and other aesthetic flaws, compared to those in good condition. Over time, a painting's varnish layer can accumulate dirt and yellow, altering its appearance. Removing the varnish can restore the original hues, but physical restoration poses risks, especially for delicate or valuable works. Traditional painting restoration uses techniques like solvent cleaning and varnish removal, though some methods may not be fully reversible and could permanently affect the artwork. Virtual restoration offers a safer alternative, allowing experimentation without damaging the original piece. Deep Learning techniques, like CNNs (J.-Y. Zhu et al., 2017) and GANs (Kumar & Gupta, 2024), have shown significant potential in image-processing tasks, including virtual art restoration. Virtual restoration involves using digital technology and software to restore and preserve artwork in a virtual environment. This process employs various techniques to repair and enhance the visual appearance of paintings without physically altering the originals. The virtual restoration method enables non-destructive testing and reversible adjustments, offering a safer approach for experimentation. Previous works used CNNs to restore yellow-filtered images to their original colors (Maali Amiri & Messinger, 2023). The research focuses on color restoration because it is one of a painting's most visually and artistically significant aspects. Accurate color restoration ensures the painting's aesthetic is preserved and the artist's original vision is maintained.

Building on prior GAN model (Wu et al., 2021), we propose using CycleGAN, which can handle unpaired data, to clean the varnish layer virtually and restore the artwork's colors. CycleGAN automatically learns to translate images between two collections, enabling effective virtual restoration. In this work, we attempt to achieve virtual restoration of artworks, specifically paintings, using deep learning, with CycleGAN as the network architecture and enhance the appearance of varnished paintings that have undergone yellowish degradation. The underlying near-original colors can be revealed by virtually removing the varnish, mitigating the yellowing effect, and restoring the artwork's authentic color tones. CycleGAN is particularly effective for tasks where paired training data is scarce or unavailable, often in art restoration. Cycle Consistency Loss is a unique feature of CycleGAN that plays a crucial role in ensuring the transformation process is reversible, thereby maintaining the integrity of the original image. This loss function helps the model learn to translate an image from one domain to another and back again, ensuring that important features and details are preserved throughout the translation process. CycleGAN has been successfully applied to style transfer tasks, which are closely related to the challenges faced in virtual painting restoration.

### 1.1 Related Work

Virtual restoration offers insight into how a painting appeared, reviving its colors to align with the artist's intent. It ensures the long-term preservation and accessibility of restoration data. Virtual restoration uses digital technology to enhance and preserve artworks without physically altering them. This approach employs non-destructive, reversible techniques, making it a safer option for experimentation. Additionally, virtual restoration is accessible to a broader audience, including researchers, artists, and enthusiasts, through appropriate software and resources (Pietroni & Ferdani, 2021).

Integrating Machine Learning in art restoration has transformed the field, enabling automated, efficient, and high-quality processes. Several approaches have been proposed: Farajzadeh & Hashemzadeh (2021) used U-Net and CNN for digital inpainting of damaged Persian pottery, focusing on noise removal, brightness adjustment, and sharpening. Sizyakin et al. (2022) employed U-Net to detect cracks in murals and GAN for inpainting the cracks. And J. Wang et al. (2023) introduced a gradient-guided dual-branch GAN to generate high-quality relic sketches. These methods highlight the versatility of machine learning techniques in addressing various restoration challenges, although they often require paired data or specific conditions for optimal performance.



H.-L. Wang et al. (2018) propose a systematic restoration framework for high-resolution deteriorated mural textures that is both efficient and effective. By using patches cropped from the original texture as training data, they achieved successful restoration results, as demonstrated in the Dunhuang mural restoration of the 61st cave. Similarly, Maali Amiri & Messinger (2021) achieved a more accurate and generalizable virtual cleaning method for paintings by using a convolutional neural network (CNN). Their approach outperforms existing physical models in restoring color quality and spectral similarity, successfully applying to famous artworks like the Mona Lisa without prior detailed information. Zou et al. developed a virtual restoration method for weathered paintings on ancient Chinese buildings using multiple deep-learning algorithms. Their approach segments the painting into the background, golden edges, and dragon patterns, applying different restoration techniques. The result provides a layered restoration that aids traditional restorers in visualizing the artwork's original appearance, reducing repetitive work and complexity (Zou et al., 2021).

J.-Y. Zhu et al. (2017) introduced an approach for image-to-image translation without the need for paired training data. Their model, CycleGAN, learns to map images from one domain to another using adversarial loss while ensuring cycle consistency through inverse mappings. This approach was tested on tasks like style transfer and photo enhancement, showing qualitative and quantitative improvements over previous methods. Engin et al. (2018) enhanced this concept with Cycle-Dehaze, an end-to-end network for single-image dehazing that combines cycle consistency and perceptual losses to improve texture recovery. Xiao et al. (2019) presented a CycleGAN-based colorization method for single grayscale images, introducing high-level semantic identity loss and low-level color loss for better optimization.

Wan et al. (2020) introduced a deep-learning method to restore old photos with severe degradation. Traditional supervised learning approaches struggle due to the complex degradation in real photos and the domain gap between synthetic and real images. To overcome this, they developed a triplet domain translation network using real photos and synthetic image pairs. They bridge the domain gap by training two VAEs to map old and clean photos into latent spaces. Their method also incorporates a global branch for structured defects (e.g., scratches) and a local branch for unstructured defects (e.g., noise), leading to superior restoration performance compared to state-of-the-art methods.

Despite these advancements in virtual restoration and machine learning applications, a research gap remains in developing models that can effectively address the unique challenges of virtual painting restoration, particularly in terms of color fidelity and detail preservation. While many existing methods focus on specific aspects of restoration or rely on paired data, there is a need for a more comprehensive approach that combines the strengths of various techniques. This research aims to fill these gaps by leveraging CycleGAN's capabilities in unpaired image translation and cycle consistency to enhance color restoration in paintings. The objectives of this study are to evaluate CycleGAN's effectiveness in maintaining color fidelity and to explore its potential for broader application in virtual art restoration.

## 2. METHODS

### 2.1 Deep Learning in Image Processing




Figure 1 Concept CycleGAN



CycleGAN (Cycle-Consistent Generative Adversarial Network) is a generative adversarial network (GAN) designed explicitly for image-to-image translation tasks where paired training data is unavailable. Introduced by J.-Y. Zhu et al. (2017), CycleGAN learns to translate images between domains without needing paired examples, making it particularly effective for applications such as style transfer, object transfiguration, and domain adaptation.

Two key data augmentation techniques were employed to effectively prepare the dataset for training the CycleGAN model to enhance its diversity and realism: yellow filtering for varnished paintings and image resizing. Yellow filtering was applied to simulate the aging effect commonly seen in varnished artworks, resulting in a yellowish or sepia-toned hue for the images.

The learning rate is a critical hyperparameter that determines the step size at each iteration toward minimizing the loss function. In CycleGAN models, a common learning rate is 0.0002, which helps ensure stable training. Learning rate decay is often employed, gradually reducing the learning rate after a specified number of epochs, such as a linear decay to zero over the final 100 epochs. This gradual reduction allows the model to converge more smoothly by making smaller weight updates as training progresses.

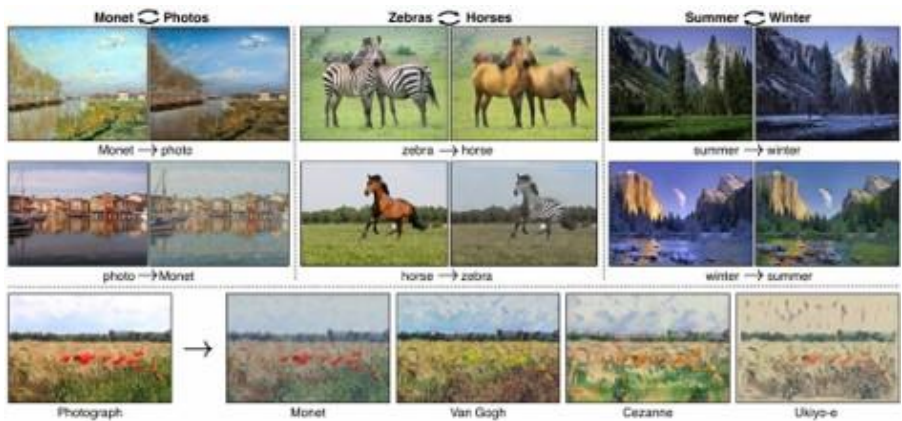


Figure 2 Example of Application CycleGAN (Maali Amiri & Messinger, 2023)

The batch size refers to the number of training examples used in a single iteration. For CycleGANs, smaller batch sizes, such as 1 or 2, are often employed. These smaller sizes are preferred because they fit better in GPU memory and are beneficial for training GANs, which can be sensitive to mini-batch statistics. Additionally, smaller batch sizes can contribute to more stable training dynamics.

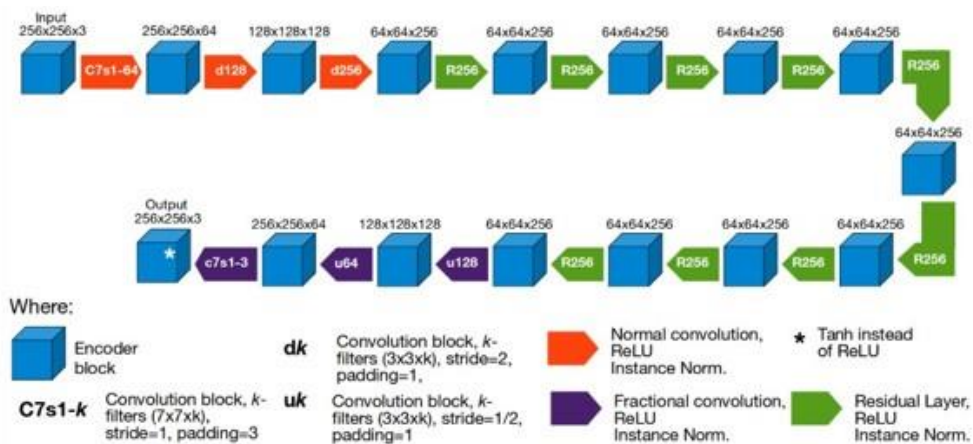


Figure 3 Internal Architecture CycleGAN



Color calibration involves adjusting and correcting the colors of an image or display to ensure they align with a standard or desired outcome. This process guarantees that the colors in the final image are accurate and consistent across different devices and media. The primary goal of color calibration is to achieve true-to-life color reproduction, minimizing discrepancies due to variations in devices, lighting conditions, and environmental factors.

Incorporating Macbeth palette images into the training dataset enhances the color calibration by providing a standardized reference that improves color accuracy. The Macbeth ColorChecker palette, featuring 24 color patches representing natural objects, is a reliable standard for adjusting colors in the restored images. This ensures that the output images remain consistent with known color values, thereby preserving the original hues and tones of the paintings (Maali Amiri & Messinger, 2023).



**Figure 4 Macbeth Color-Checker Palette (1 – 4 filtered as varnished reference, Color-Checker as original/unvarnished reference)**

## 2.2 Experimental Dataset and Parameter Environment

For this study, the dataset used in the virtual restoration process was collected from several reputable online sources. These websites provided high-quality digital images of various artworks, which were selected based on specific criteria such as degradation, varnish yellowing, and visual impairments commonly addressed in restoration. The collected dataset was then preprocessed and curated to ensure diversity in art styles and conditions, allowing for a comprehensive evaluation of the proposed CycleGAN model. This curated dataset played a crucial role in training and validating the model for effective virtual restoration. However, to further enhance the model's reliability, it is important to provide more quantity and variety of art styles and degradation conditions within the dataset. Adding more data variety, including abstract styles or other types of damage, would help strengthen the model's performance. Additionally, plans or suggestions for increasing data variety in future work should be considered, as they would contribute to a more robust and versatile model capable of handling a more comprehensive range of restoration challenges.

Several techniques augment the dataset, such as random cropping, which extracts patches from images to increase diversity and reduce overfitting. Scaling and resizing images to a fixed size (256x256 pixels) standardizes input dimensions for efficient training. Additionally, techniques like horizontal and vertical flipping and color jittering (random adjustments to brightness, contrast, saturation, and hue) introduce variations for the model to learn.






Figure 5 The result of the resize step

### 2.3 Testing and Performance Metrics

PSNR (Peak Signal-to-Noise Ratio) offers a clear numerical representation of the difference between the restored and reference images. It is a valuable metric for assessing the overall quality and fidelity of the restoration process, as higher PSNR values indicate more accurate restorations with less distortion or noise.

SSIM (Structural Similarity Index Measure) assesses the perceptual quality of images by considering changes in structural information essential for human visual perception. This makes SSIM especially relevant for artistic image restoration tasks, where preserving structural integrity is vital for a faithful representation of the original artwork.

PSNR quantifies the ratio between the maximum possible power of a signal (image) and the power of corrupting noise that degrades its fidelity. This metric is expressed in decibels (dB) and is commonly used to assess the overall quality of image restoration processes.

$$PSNR = 10 \cdot \log_{10} \left( \frac{MAX^2}{MSE} \right) \quad (1)$$

MAX represents the image's maximum pixel value (typically 255 for 8-bit images). At the same time, MSE (Mean Squared Error) is the average of the squared differences between corresponding pixels of the restored and reference images.

A higher PSNR value indicates better image quality, suggesting lower distortion or noise in the restored image compared to the reference image. Because of its ability to quantify the fidelity of reconstructed images, PSNR is widely used in image restoration tasks.

Euclidean distance measures the straight-line distance between two points in multi-dimensional space. In image restoration, it quantifies the pixel-wise differences between the restored and reference images, indicating how closely the restoration matches the original.

### 2.4 Implementation

Training CycleGAN poses several challenges, with mode collapse being a significant issue, where the generator produces limited output diversity. This can be mitigated by adding noise to inputs or using different mini-batches to update the generators and discriminators. Balancing adversarial and cycle consistency losses is crucial for stability, and adjusting the weights  $\lambda_{cyc}$  and  $\lambda_{identity}$  can aid in achieving this balance. Techniques such as gradient clipping can help prevent



exploding gradients, and using instance normalization instead of batch normalization enhances training stability and performance. Finally, due to the resource-intensive nature of CycleGAN training, employing strategies like mixed-precision training and distributed training across multiple GPUs can significantly accelerate the process.

In this work, the CycleGAN model was configured with a batch size of 1, a learning rate 0.0002, and cycle loss lambda values of 10. The model was trained for 1000 epochs, with regular checkpoints saved to monitor progress and prevent overfitting.

The batch size is set to 1, allowing the model to update its parameters based on each image pair's gradient information. This configuration helps manage memory constraints while facilitating fine-grained updates. Additionally, with Cycle Consistency Loss set to LAMBDA\_CYCLE = 10, the model ensures consistency in transformations between varnished and unvarnished images.

CycleGAN's optimization employs Adam optimizers for the generators and discriminators, leveraging its effectiveness with large datasets and an adaptive learning rate mechanism. Typical parameters for the Adam optimizer include  $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$ , and a learning rate of 0.0002, striking a balance between convergence speed and stability. Additionally, a linear learning rate decay is applied after a specified number of epochs (e.g., after the first 100 epochs in a 200-epoch training cycle), aiding in fine-tuning the model towards the end of training.

### 3. RESULTS AND DISCUSSION

#### 3.1 Qualitative Results

The results showcase the visual results of the CycleGAN model applied to the restoration of varnished paintings, demonstrating its effectiveness in removing and restoring the original appearance. Notable examples include Monet's "Nasturtiums in a Blue Vase" and Hermann Corrodi's "Prayer at Dawn," which feature different styles and restoration challenges, highlighting the model's robustness and versatility.



Figure 6 Original Unvarnish Traditionally (left), Virtual Unvarnish Restoration (middle), Original Varnished Painting (right)




Figure 7 Original Unvarnish Traditionally (left), Virtual Unvarnish Restoration (middle), Original Varnished Painting (right)



Figure 6 features a still life by Claude Monet, where the varnish has dulled the vibrancy of the flowers and background. The original varnished image displays a significant yellow tint that obscures the vivid colors of the nasturtiums and blue vase. In contrast, the restored image reveals Monet's intended true colors, showcasing the flowers' vibrancy and the blue vase's prominence.

Figure 7 features a historical artwork by Hermann Corrodi depicting a serene dawn prayer scene, affected by varnish that has dimmed its overall brightness. The original image appears muted, with a yellowish haze reducing the clarity and vibrancy of the scene. In contrast, the restored image reveals richer, more accurate colors, enhancing the prominence and clarity of the dawn light and architectural details.

This comparative analysis underscores the strengths of the CycleGAN model in virtual painting restoration. While the lack of direct visual comparisons is limited, the model's quantitative performance and qualitative improvements are promising. Future work should focus on comprehensive visual comparisons and expanding datasets to validate our findings further.

The comparison of paintings before and after virtual restoration reveals significant improvements in several key areas. One of the most noticeable enhancements is the improved color accuracy. The restoration model effectively eliminates the yellow tint caused by varnish, thereby restoring the original hues intended by the artist. This correction not only revives the aesthetic appeal of the artwork but also ensures that the artist's original vision is preserved. In addition to color restoration, the model demonstrates proficiency in recovering fine details previously obscured by the varnish. This capability is crucial for maintaining the integrity and authenticity of the artwork, as it allows viewers to appreciate the intricate details and textures that define the painting.

Furthermore, the restoration often produces a cleaner image with reduced noise, offering a clearer artwork view. This noise reduction enhances the overall visual quality and allows a more accurate interpretation of the painting's features and elements. Together, these improvements underscore the effectiveness of the restoration model in enhancing both the aesthetic and historical value of artworks.

### 3.2 Quantitative Results

This section evaluates the performance of the CycleGAN model using various quantitative metrics, which provide a numerical assessment of the model's ability to restore varnished paintings. The metrics employed include Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), Euclidean Distance, and Color Inspector 3D. These metrics collectively offer insights into the model's effectiveness in enhancing image quality, preserving structural integrity, and accurately restoring colors. By analyzing these metrics, we can objectively assess the CycleGAN model's performance in virtual painting restoration.

PSNR (Peak Signal-to-Noise Ratio) measures the ratio of the maximum possible signal power to the power of corrupting noise, quantifying image quality by comparing the original unvarnished image to the restored image. SSIM (Structural Similarity Index Measure) evaluates image quality based on structural information, luminance, and contrast between the original and restored images, with values ranging from -1 to 1; higher values indicate greater structural similarity.

Euclidean Distance calculates the straight-line distance between corresponding pixels in the RGB color space of the original and restored images, with lower distances indicating higher color fidelity. The 3D Color Inspector provides a visual representation of the color distribution in the original and restored images within 3D RGB space, allowing for visual comparison to identify discrepancies.

Several aspects are considered to evaluate the quality of the results. First, the Average Euclidean Distance is assessed, where a low value suggests that the two images are very similar, while a high value indicates significant color differences. Second, the Distance Distribution is analyzed; a graph with most points near the Y-axis, indicating small distance values, suggests that many





pixels are very similar in both images, whereas a wide distribution implies large variations, indicating potential differences between the images. Third, Peaks in the Graph are examined; high peaks at low distance values indicate that many pixels are almost identical in the two images, while scattered peaks suggest variations in pixel similarity. Finally, PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index Measure) are considered, with a high PSNR typically indicating very similar images and an SSIM value close to 1, signifying a high degree of structural similarity. These aspects collectively provide a comprehensive assessment of image quality and similarity.

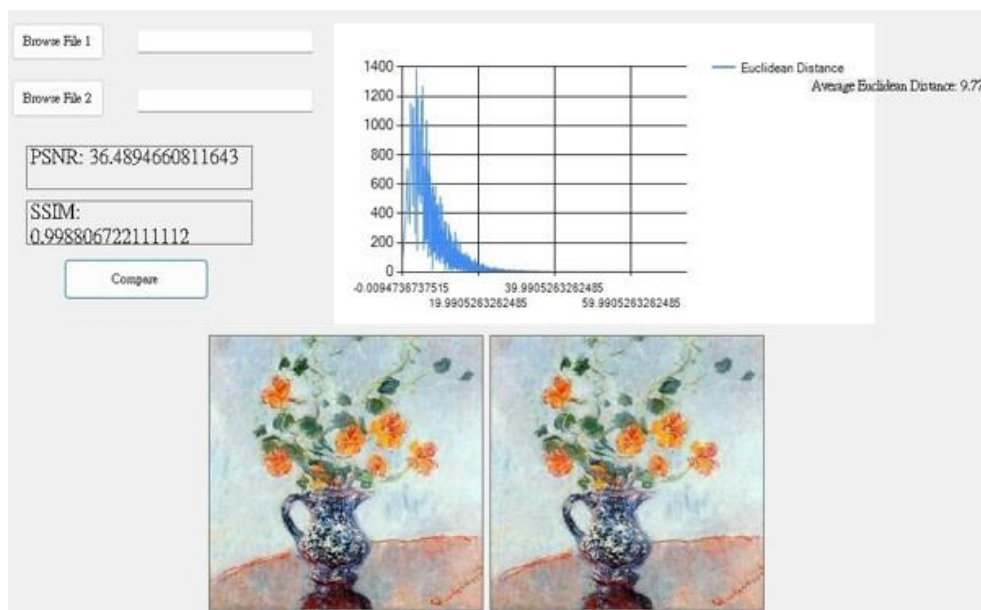


Figure 8 Result of Monet painting analysis PSNR, SSIM, and Euclidean Distance

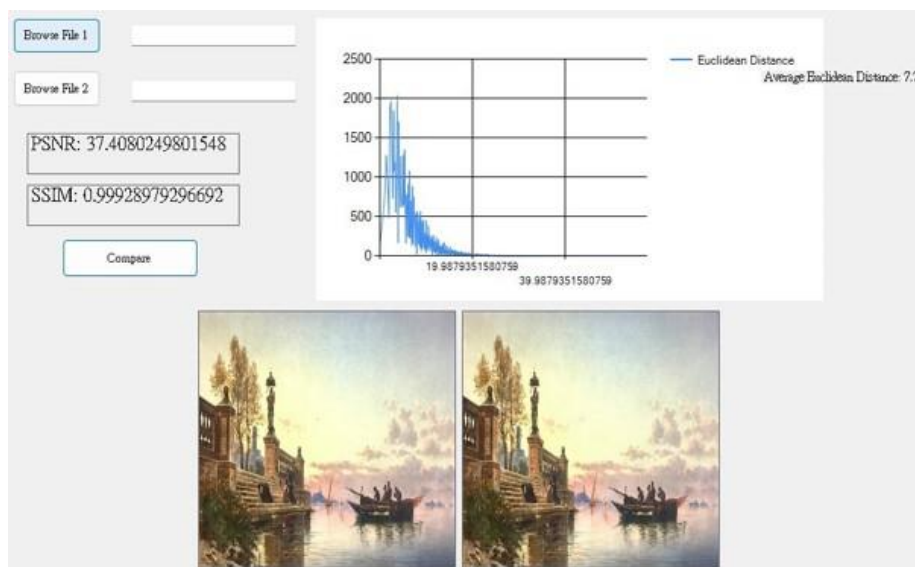


Figure 9 The result of Hermann painting analysis is PSNR, SSIM, and Euclidean Distance

Understanding the example results involves analyzing image similarity through specific metrics. As shown in Figure 8 and Figure 9, highly similar images are indicated by points on the graph that are close to the Y-axis, reflecting a high Peak Signal-to-Noise Ratio (PSNR), typically above 30



dB, and a high Structural Similarity Index Measure (SSIM), often above 0.9. These metrics suggest superior image quality and strong similarity between images. Conversely, dissimilar images exhibit a wide distribution of points on the graph, with low PSNR values, typically below 20 dB, and low SSIM scores, often below 0.5. PSNR values serve as image quality indicators, with higher values denoting better quality. SSIM assesses the similarity between two images based on luminance, contrast, and structure, ranging from -1 to 1, where 1 signifies perfect similarity. These metrics collectively provide a comprehensive understanding of image quality and similarity.

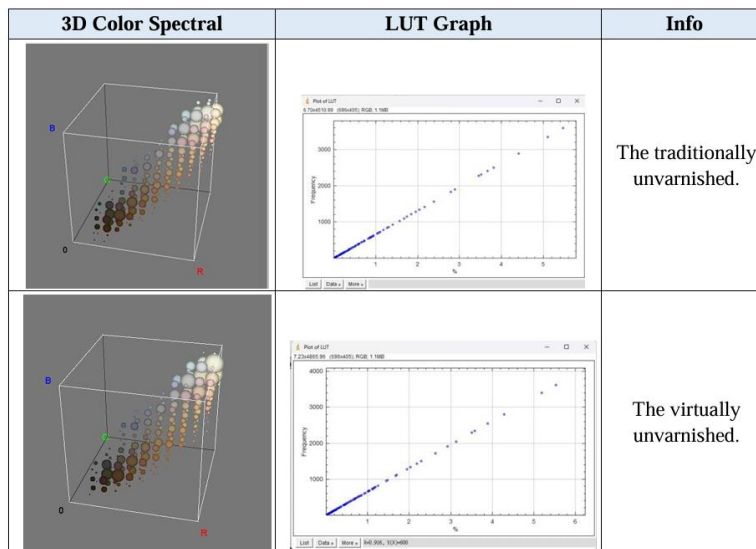


Figure 10 Corrodi - "Prayer at Dawn" Color Inspector 3D Simulation

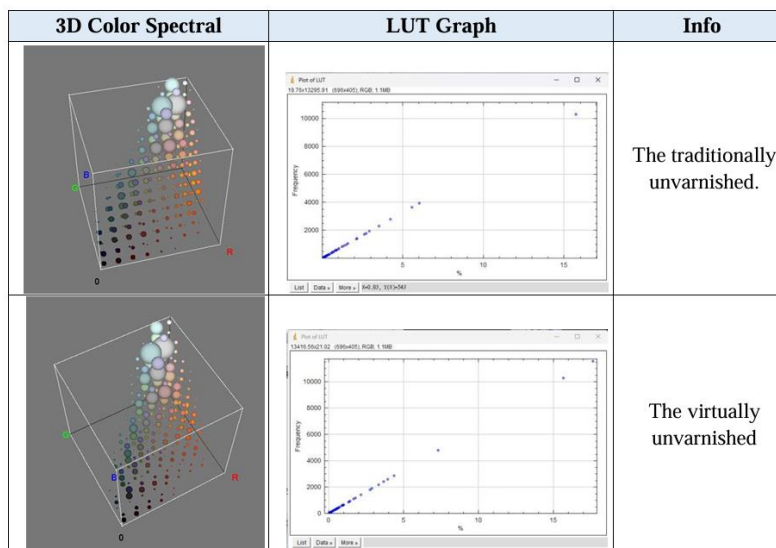


Figure 11 Monet - "Nasturtiums in a Blue Vase" Color Inspector 3D Simulation

Figures 10 and 11 utilize Color Inspector 3D to visually and quantitatively evaluate the quality of color restoration. Consistent color distributions in these figures indicate successful restoration efforts. The Look-Up Table (LUT) graph is particularly useful, as it displays how colors in the input image are mapped to those in the output (restored) image. By including the LUT graph for this painting, specific observations can be discussed, such as areas where color consistency is well maintained and any noticeable shifts that may occur. Additionally, the 3D Cube RGB spectral



color feature in Color Inspector 3D is an advanced tool for visualizing and analyzing the distribution of colors within an image in a three-dimensional space. This involves separating the image into Red, Green, and Blue channels to study each color's contribution to the overall image. By comparing the 3D RGB cubes of the input (varnished) and output (restored) images, we can assess how the color distribution has changed due to the restoration process. This comparison can reveal whether the model is accurately restoring the original colors or introducing any color distortions.

A well-restored image should have a color distribution closely matching the original unvarnished painting. By comparing the 3D RGB cubes of the input and output images, it can visually demonstrate the accuracy of the color restoration, as shown in Figures 10 and 11.

**Table 1 Comparison with Other Works**

No.	Method	PNSR (dB)	SSIM	Source
1	GAN – Artwork Restoration	28.90	0.53	Kumar & Gupta (2024)
2	VAE (Variational Autoencoder) – Old Photo Restoration	23.33	0.70	Wan et al. (2020)
3	Conditional GAN – Artwork Inpainting Restoration	N/A	0.795	Adhikary et al. (2021)
4	CycleGAN – Image Dehazing	15.54	0.66	Engin et al. (2018)
5	CNN – Based Inpainting	22.14	N/A	Zeng et al. (2020)
6	GAN – Mural Restoration	34.36	0.91	Li et al. (2021)
7	CycleGAN – Virtual Unvarnished Painting Restoration	36.95	0.998	<i>current work</i>

Table 1 compares various deep-learning approaches applied to color enhancement, image restoration, and artwork restoration. However, it is essential to note that this comparison does not utilize the same dataset across all models, which introduces certain limitations. Differences in datasets, preprocessing steps, and training parameters can significantly impact the comparability of the results. These variations may lead to discrepancies in performance metrics and outcomes, making it challenging to draw definitive conclusions about the relative effectiveness of each approach. Therefore, while the table offers valuable insights into the capabilities of different models, caution should be exercised when interpreting the results due to these inherent differences.

#### 4. CONCLUSIONS

Comparing the CycleGAN model with traditional painting restoration methods and other state-of-the-art techniques helps to position the work within the broader field. This comparison highlights the strengths and weaknesses of our approach, providing valuable insights for future improvements.

Traditional painting restoration methods rely heavily on manual techniques performed by skilled restorers. While these methods effectively preserve and restore artworks, they are often time-consuming and susceptible to human error. Key issues include maintaining consistency in color restoration and the potential for human error, which can lead to unintended alterations in the artwork. These challenges underscore the need for more efficient and reliable restoration techniques that complement or enhance traditional methods. Virtual restoration is a complementary tool to the physical restoration of artwork, offering insights into potential results without replacing the traditional restoration process. It gives restorers a preview of how a painting might look post-restoration, assisting in decision-making and planning.

The CycleGAN model excelled in color restoration and detail preservation, as indicated by higher PSNR and SSIM scores than other methods. This success can be attributed to the cycle



consistency loss, which ensures that the restored images remain faithful to their original, unvarnished versions, thereby maintaining the structural integrity critical for preserving artistic details and nuances.

Expanding the dataset is crucial for improving the model's ability to handle diverse styles and conditions, as the limited diversity in the current training dataset may lead to artifacts, especially when dealing with unique paintings like abstract artworks. Plans to explore additional datasets, including those featuring various art styles and degradation types, can provide a more comprehensive training ground for the model. The absence of similar examples in the training data can result in poor performance and undesirable artifacts due to the variable nature of abstract art, disrupting the training process and yielding inconsistent outputs. Additionally, experimenting with other models like StyleGAN or VQ-VAE could offer new perspectives and techniques in art restoration. These models, known for their high-quality generative capabilities, might enhance the restoration process by providing more nuanced and style-specific restorations, thus addressing the challenges posed by limited dataset diversity.

The effectiveness of cycle consistency loss ensures bidirectional and stable transformations between varnished and unvarnished images. Additionally, specific data augmentation strategies—such as yellow filtering for varnished paintings and resizing techniques—improved the model's ability to generalize across different painting styles, resulting in a more robust training dataset that enhanced overall performance. Unique preprocessing steps, including yellow varnish filtering and standardizing image sizes, further tailored the dataset for the restoration task, boosting the model's effectiveness.

In conclusion, our study effectively restored the color fidelity and nuances of historical paintings using CycleGAN. This model is a complementary tool to physical restoration, offering valuable insights without replacing traditional methods. Our findings highlight CycleGAN's potential in digital art conservation and underscore the importance of ongoing research and interdisciplinary collaboration.

When using CycleGAN for restoration, several limitations must be considered, particularly with paintings with additional layers or reconstructions obscuring the original artwork. New layers of paint can conceal original details or alter the original composition, making it challenging for the CycleGAN model to accurately predict details hidden beneath these overpainted layers.

This research paves the way for further exploration in enhancing virtual restoration techniques, potentially integrating more advanced deep learning models and expanding datasets to include a wider range of artistic styles and conditions. By bridging the gap between traditional conservation practices and cutting-edge technology, CycleGAN provides art conservators with a valuable tool for informed decision-making and the preservation of cultural heritage.

Interdisciplinary collaboration is crucial in digital restoration research, as working with experts in computer vision, art history, and materials science can yield new insights and perspectives. These cross-disciplinary efforts can lead to more comprehensive restoration techniques that address the technical and artistic aspects of virtual painting restoration. Partnering with art historians and materials scientists enhances understanding of artworks' historical, aesthetic, and physical contexts, while collaboration with art conservation professionals ensures alignment with traditional practices. Overall, CycleGAN exemplifies how technology can transform art restoration, merging art and technology to preserve and revitalize paintings, thereby protecting cultural heritage for the future.

## REFERENCES

- Adhikary, A., Bhandari, N., Markou, E., & Sachan, S. (2021). ArtGAN: Artwork Restoration using Generative Adversarial Networks. *2021 13th International Conference on Advanced Computational Intelligence (ICACI)*, 199–206. <https://doi.org/10.1109/ICACI52617.2021.9435888>



- Engin, D., Genc, A., & Ekenel, H. K. (2018). Cycle-Dehaze: Enhanced CycleGAN for Single Image Dehazing. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 938–9388. <https://doi.org/10.1109/CVPRW.2018.00127>
- Farajzadeh, N., & Hashemzadeh, M. (2021). A deep neural network based framework for restoring the damaged persian pottery via digital inpainting. *Journal of Computational Science*, 56, 101486. <https://doi.org/10.1016/j.jocs.2021.101486>
- Kumar, P., & Gupta, V. (2024). Unpaired Image-to-Image Translation Based Artwork Restoration Using Generative Adversarial Networks. In *Smart Innovation, Systems and Technologies* (Vol. 372, pp. 581–591). [https://doi.org/10.1007/978-981-99-6774-2\\_52](https://doi.org/10.1007/978-981-99-6774-2_52)
- Li, J., Wang, H., Deng, Z., Pan, M., & Chen, H. (2021). Restoration of non-structural damaged murals in Shenzhen Bao'an based on a generator–discriminator network. *Heritage Science*, 9(1), 6. <https://doi.org/10.1186/s40494-020-00478-w>
- Maali Amiri, M., & Messinger, D. W. (2021). Virtual cleaning of works of art using deep convolutional neural networks. *Heritage Science*, 9(1), 94. <https://doi.org/10.1186/s40494-021-00567-4>
- Maali Amiri, M., & Messinger, D. W. (2023). Virtual cleaning of works of art using a deep generative network: spectral reflectance estimation. *Heritage Science*, 11(1), 16. <https://doi.org/10.1186/s40494-023-00859-x>
- Pietroni, E., & Ferdani, D. (2021). Virtual Restoration and Virtual Reconstruction in Cultural Heritage: Terminology, Methodologies, Visual Representation Techniques and Cognitive Models. *Information*, 12(4), 167. <https://doi.org/10.3390/info12040167>
- Sizyakin, R., Voronin, V. V., & Pizurica, A. (2022). Virtual restoration of paintings based on deep learning. In W. Osten, D. Nikolaev, & J. Zhou (Eds.), *Fourteenth International Conference on Machine Vision (ICMV 2021)* (p. 60). SPIE. <https://doi.org/10.1117/12.2624371>
- Wan, Z., Zhang, B., Chen, D., Zhang, P., Chen, D., Liao, J., & Wen, F. (2020). *Bringing Old Photos Back to Life*. <http://arxiv.org/abs/2004.09484>
- Wang, H.-L., Han, P.-H., Chen, Y.-M., Chen, K.-W., Lin, X., Lee, M.-S., & Hung, Y.-P. (2018). Dunhuang mural restoration using deep learning. *SIGGRAPH Asia 2018 Technical Briefs*, 1–4. <https://doi.org/10.1145/3283254.3283263>
- Wang, J., Zhang, E., Cui, S., Wang, J., Zhang, Q., Fan, J., & Peng, J. (2023). GGD-GAN: Gradient-Guided dual-Branch adversarial networks for relic sketch generation. *Pattern Recognition*, 141, 109586. <https://doi.org/10.1016/j.patcog.2023.109586>
- Wu, Y., Wang, X., Li, Y., Zhang, H., Zhao, X., & Shan, Y. (2021). Towards Vivid and Diverse Image Colorization with Generative Color Prior. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 14357–14366. <https://doi.org/10.1109/ICCV48922.2021.01411>
- Xiao, Y., Jiang, A., Liu, C., & Wang, M. (2019). Single Image Colorization Via Modified CycleGAN. *2019 IEEE International Conference on Image Processing (ICIP)*, 3247–3251. <https://doi.org/10.1109/ICIP.2019.8803677>
- Zeng, Y., Gong, Y., & Zeng, X. (2020). Controllable digital restoration of ancient paintings using convolutional neural network and nearest neighbor. *Pattern Recognition Letters*, 133, 158–164. <https://doi.org/10.1016/j.patrec.2020.02.033>
- Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). *Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks*. <https://doi.org/https://doi.org/10.48550/arXiv.1703.10593>
- Zou, Z., Zhao, P., & Zhao, X. (2021). Virtual restoration of the colored paintings on weathered beams in the Forbidden City using multiple deep learning algorithms. *Advanced Engineering Informatics*, 50, 101421. <https://doi.org/10.1016/j.aei.2021.101421>

