

# Scaling Art Classification Models: Enhancing Binary Classifiers and Tackling the Challenge of AI-Generated Art

Brendan Gorman

*Brendan Gorman is projected to graduate in Spring 2025 with a major in Computer Science. His research interests are related to computer vision, AI, cybersecurity, and machine learning.*

**Abstract:** This paper discusses the creation of an AI-based binary classification model that can efficiently classify artwork by artists, with two categories building on previous research in the area. The research topic here is to streamline current multi-classification models by using a binary classifier in assessing its performance compared to conventional multi-classification systems with multiple categories and artists simultaneously. Deep learning methods, specifically ResNet-101, were employed to distinguish between non-Monet and Monet paintings in the first study, and between Vincent van Gogh and non-van Gogh paintings in the second. The paper also discusses the implications of Artificial Intelligence (AI)-generated pieces of art, briefly delving into the difficulty of identifying if artworks are genuine or not to a computing system. The results show that it is possible to develop a theoretical multi-classifier through the fusion of various binary classifiers, thereby creating an efficient and scalable approach for handling large datasets and many artists. Nonetheless, whereas binary classification proves to be effective for traditional art with respect to accuracy, it cannot differentiate AI paintings imitating artists, thereby representing the limitation of the method. The paper concludes by emphasizing the potential for further advancements in art classification, particularly considering the growing impact of AI-generated artworks.

**Research Advisors:** Todd Dobbs, College of Computing and Informatics

**Key Words:** Binary classification, AI-generated art, art authentication, ResNet-101, scalability, Money paintings

## Introduction

In recent years, the art world has faced unprecedented challenges in authentication and verification, particularly with the rise of sophisticated forgeries and AI-generated artwork. Thomas Hoving, former director of the Metropolitan Museum of Art, estimated that at least 50,000 pieces in circulation were inauthentic. In his book *False Impressions*, he famously stated, “I

almost believe that there are as many bogus works as genuine ones” [5]. This sentiment underscores the growing limitations of traditional authentication methods, which rely heavily on expert analysis and provenance documents. In an era where both human forgers and artificial intelligence can create remarkably convincing replicas, these conventional approaches are no longer sufficient. As a result, museums, collectors, and researchers

are increasingly turning to scientific techniques and machine learning models to enhance the accuracy of authentication.

In light of these developments, deep learning classification systems have arisen as a potentially worthwhile response to the need for art authentication. While stable analytical capabilities are provided by such models, they also present profound challenges. Recent research indicates that distributed deep learning models suffer from serious issues of computational complexity as well as scalability, especially in multi-graphical user interface setups. Evaluation results indicate considerable variation in the scalability of these frameworks, and the need for load balancing for parallel distributed deep learning [6]. To put it simply, these models require significant computational resources, and their efficiency varies depending on how well the workload is distributed across multiple GPUs. Deep learning frameworks such as TensorFlow, MXNet, and Chainer offer numerous primitive elements required for effective neural network structure creation for a variety of applications, such as computer vision, speech recognition, and natural language processing. These software tools are essential to help build and train models by providing functions for data processing, optimization, and efficient computation. Nonetheless, the performance of the deep learning frameworks at execution time is considerably inconsistent even when training the same deep network models on the same GPUs [6]. Although deep learning models such as deep convolutional neural networks (DCNNs) and ResNet have been promising for art classification, they need a lot of fine-tuning and computational power to attain high accuracy. The pre-trained DCNN demonstrates how fine-tuning DCNNs on large-scale artistic collections considerably enhances classification

performance, enabling the networks to learn new selective attention mechanisms over the images [6]. Furthermore, this demonstrates that DCNNs that have been fine-tuned on a large artwork dataset classify cultural heritage objects from another dataset better than the same architectures pre-trained solely on the ImageNet dataset [9].

This research proposes a new approach using binary classification frameworks to overcome such challenges. By focusing exclusively on binary classification, it will seek to build a more efficient and less complex system that can seamlessly adapt to a range of artists while maintaining high accuracy. This work extends current work in artwork authentication [2] and introduces new work in model architectures and training experiments involving AI artworks.

## **Motivations and Research Objectives**

The increasing prevalence of artwork generated through artificial intelligence, together with high-profile cases of artwork forgery, have underpinned the need for sophisticated methodologies for authenticating artwork. In 2018, the J. Paul Getty Museum in Los Angeles finally acknowledged that its ancient Greek sculpture, the Getty Kouros, was in fact a contemporary forgery. The museum acquired the sculpture in 1985 for approximately \$9 million; however, subsequent investigations demonstrated the statue's atypical and chronologically incompatible shape, along with fake provenance documents, indicate it was a forgery made recently. [1]. This case exemplifies how vulnerabilities in provenance records can be exploited and highlights the necessity for more robust verification techniques in both traditional and AI-generated art. With growing availability of software to create AI artwork, it has become increasingly challenging to discern

between artwork produced through human hands and artwork generated through AI. For instance, in 2022, a digital painting “Théâtre D’opéra Spatial” was created by Jason M. Allen through an AI website called Midjourney and had won a digital artwork competition in Colorado, sparking controversy over whether AI can be used to create valid artwork [10]. Moreover, Allen also encountered problems related to the unauthorized use and merchandising of his work [11]. Collectively, these developments highlight the growing demand for sophisticated AI-powered systems for artwork authentication.

To address these concerns, this research study identifies three key objectives. First, it aims to develop a binary classification model capable of distinguishing between works of art produced by Claude Monet and works produced by other artists, providing a reliable tool for confirming the legitimacy of Monet’s style. Second, generalizability of the model is evaluated through its application to works produced by other renowned painters, such as Vincent van Gogh, in an examination of the model’s capability to classify a variety of forms with little variation in its output. Third, the effectiveness of the model will be evaluated by its ability to differentiate between works produced by humans and works produced through artificial intelligence, testing for the potential for AI-created replicas of Monet’s style to deceive the classifier. Together, these objectives stand to expand our understanding of artificial intelligence in relation to artwork analysis and lay a basis for future development in terms of proper use of AI in artwork authentication, with a view towards resolving growing concerns regarding artwork forgery and integrity in virtual artwork markets.

## **Methodology**

This study utilizes binary classification, and its use is supported by a variety of key factors. Earlier experiments in single-class classification failed, with the models tending to overfit and classify most works of art as Monet, therefore not having acquired meaningful artwork features. In addition, multi-class classification holds a lot of potential but suffers a drop in accuracy with an increased number of artists, with values dropping from 91.23% for 100 artists to 48.97% when moving to 2,368 artists [2]. As such, binary classification yields better accuracy and increased scalability for use in artwork authentication.

ResNet-101 was utilized for its proven effectiveness in resolving vanishing gradient problem through residual connections, allowing for an unobstructed flow of gradients through skip connections between deeper and preceding layers and to the first filters. In other words, the vanishing gradient problem makes it difficult for deep networks to learn effectively, but ResNet overcomes this by using shortcut connections that help preserve important information as it moves through the layers. This architectural feature has proven effective in image classification, allowing for training deeper networks with model accuracy preserved. The use of ResNet deep architecture in combination with careful use of techniques for data augmentation enables the model to learn pertinent features for distinguishing between types of artworks and reduces overfitting challenges.

## **Experimental Setup/Design**

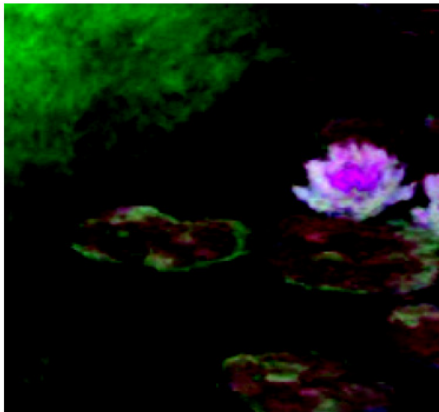
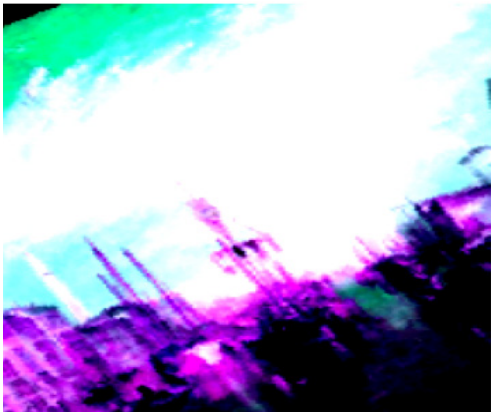
The experiment started with a thorough preprocessing pipeline with the goal of optimizing model performance. As seen in Table 1, there is a list of critical transformations for uniformity: first, a resize of images, then a rotation, with

a 50% chance of a horizontal flipping for generating orientation variance. Shear transformations were also utilized for compensating for perspective mismatches, and color jittering for allowing randomness in terms of

Table 1. Summary of image preprocessing transformations applied during training.

Transformation	Specification
Image Resize:	Crops to 224x224 pixels
Image Rotation:	Rotates between -10 to +10 degrees and a 50% probability to horizontally flip
Image Shear:	Shears from -5 to +5 degrees
Image Color Jitter:	Randomly changes brightness saturation and hue
Image Tensor Conversion:	Change the image to tensor format for model processing
Image Normalize:	Changes mean and standard deviation values for Resnet

Table 2. Two images after processing, finalized and ready for usage. The left image represents a Monet painting, while the right image represents a non-Monet painting.

Monet Image	Non-Monet Image
	

brightness, saturation, and hue. Lastly, tensor conversion and normalization specific to requirements for use with ResNet took place to finalize the image seen in Table 2.

The model's architecture then employed a pre-trained ResNet-101 that was adjusted, specifically for use in binary classification represented in Figure 1. The network begins with an RGB input image of  $224 \times 224 \times 3$ . An RGB image is a digital image where each pixel is composed of three-color channels—red, green, and blue—that combine in varying intensities to produce a wide range of colors. The model then progresses through five layers of convolution (Conv2-Conv5), designed to scan an image for a variety of features including textures, borders, and patterns, with incorporated residual phases allowing for information to pass through unscathed in its journey through the network. Next, average pooling is utilized for flattening, then followed by full connected layers and a dropout layer which randomly deactivates some of the connections to prevent overfitting, culminating in the overall output of binary classification between Monet and non-Monet prediction. This configuration enables efficient feature extraction with proper maintenance of gradient flow through the long network structure.

The model's architecture then employed a pre-trained ResNet-101 that was adjusted, specifically for use in binary classification represented in Figure 1. The network begins with an RGB input image of  $224 \times 224 \times 3$ . An RGB image is a digital image where each pixel is composed of three-color channels—red, green, and blue—that combine in varying intensities to produce a wide range of colors. The model then progresses through five layers of convolution (Conv2-Conv5), designed to scan an image for a variety of features including textures, borders, and patterns, with incorporated residual phases allowing for information to pass through unscathed in its journey through

the network. Next, average pooling is utilized for flattening, then followed by full connected layers and a dropout layer which randomly deactivates some of the connections to prevent overfitting, culminating in the overall output of binary classification between Monet and non-Monet prediction. This configuration enables efficient feature extraction with proper maintenance of gradient flow through the long network structure.

The actual training process, as can be observed in the middle of Figure 2, had two iteration blocks per epoch. The first block iterated over initial batches of paintings (Pr-1 through Pr-32), and the second block iterated over batches of larger quantities (Pr-1957 through Pr-1988). Validation after every epoch was done with a dedicated validation set (Pv-1 through Pv-552) to facilitate continuous monitoring of model performance and easy identification of possible overfitting.

The training deployment, in PyTorch, used GPU acceleration with Central Processing Unit fallback support. Performance was tracked using exhaustive metric monitoring, including training and validation loss, accuracy measures, and F1 measures for the Monet class (see Appendix D for a detailed breakdown and the complete training data document). The validation procedure, depicted on the right side of Figure 2, used a similar process to the training phase but without backpropagation. It evaluates its performance based on the unseen data using the same loss, accuracy and F1-score and it prints a confusion matrix for reference.

To avoid overfitting, the early stopping method was invoked with patience of five epochs and automatically saved the model with the best performing validation F1 scores as well as backpropagation in the training phase. This will in turn teach the network to make better predictions over time by constantly refining its weights based

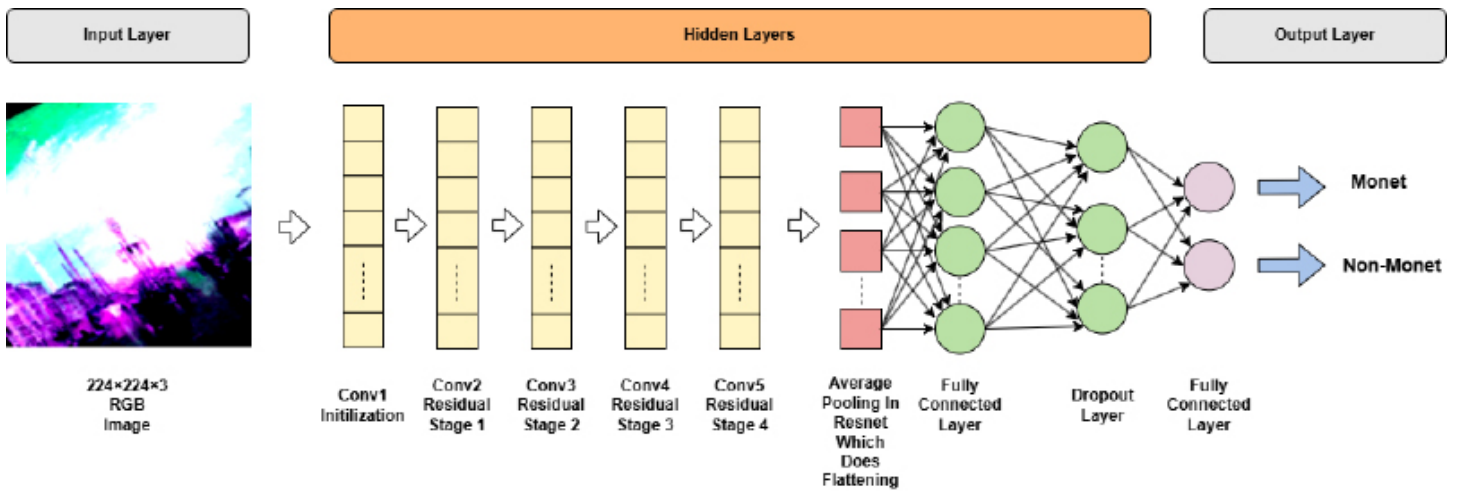


Figure 1. Monet vs non-Monet painting binary classification ResNet-101 architecture. The network takes in a  $224 \times 224 \times 3$  RGB image and performs several stages of convolutional residuals to it. The final feature maps are subjected to average pooling, followed by fully connected layers and a dropout layer before Monet vs non-Monet classification.

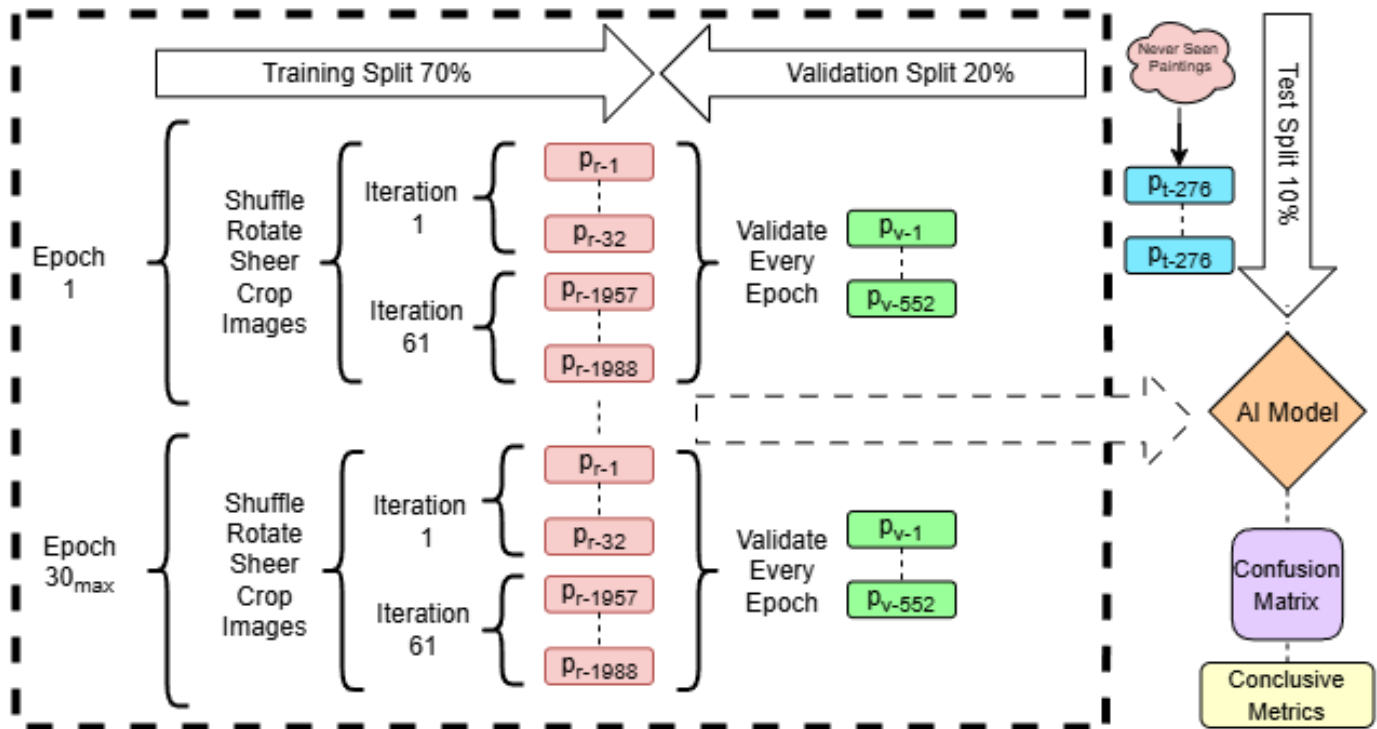


Figure 2. Mapped out distribution of datasets used for training, validation, and testing in the model. The training set is represented by  $P_r-i$ , the validation set  $P_v-i$ , and the testing set  $P_t-i$ . All data sets are mutually exclusive. Each epoch is one complete pass of the entire training dataset through the model.

on how wrong its predictions were and how to correct them accordingly helping the model learn from its mistakes by adjusting and improving for the next round. This entire process of training and testing the model guarantees reproducibility while allowing the flexibility required for tailoring the framework to be used with other artists, as outlined in the research goals.

The testing phase employed the held-out test set (Pt-276) to determine the performance of the finalized model. This testing involved the examination of the confusion matrix and absolute metrics, providing a global accounting of the model’s capacity to distinguish Monet and non-Monet works. The pipeline is a robust model for binary classification of art authenticity, with possible uses for both classical forgery detection and artificial generation detection.

**Results:**

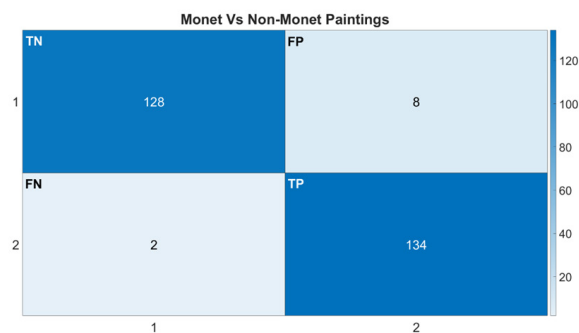
The binary model demonstrated high competence in discriminating between real Monet paintings and non-Monet paintings but with intriguing challenges when confronted with AI-generated images. The results can be examined using various significant features:

**1. Model Training and Convergence**

The training dynamics in Figure 3 display good model convergence by loss and accuracy metrics over 12 epochs. The training loss exhibited a consistent drop from a level of approximately 0.45 to 0.15, while the validation loss flattened out at a level near 0.20, reflecting effective generalization with minimal overfitting. The accuracy plots indicate an early jump in improvement, which converged to training and validation accuracy levels near approximately 93% by epoch 12. The concurrent trends of the training and validation measures



Figure 3. Training Progress of Monet AI Model: Loss and Accuracy Over 12 Epochs. The top graph shows the decrease in both training and validation loss, while the bottom graph illustrates the increase in training and validation accuracy, indicating successful learning and generalization of the model.



Accuracy	Precision	Recall	F1 Score
96.32%	94.37%	98.53%	96.4%

Figure 4. Finalized results from the Monet Model comparing the Monet and Non-Monet datasets, including the confusion matrix and performance metrics.


reflect an evenly paced learning process that successfully skirted the usual traps of overfitting or underfitting.

## 2. Classification Performance on Traditional Artworks

The model was very successful at differentiating between Monet paintings and non-Monet paintings, as presented in the confusion matrix in Figure 4. The performance is 128 true negatives (correct identification of non-Monet paintings) and 134 true positives (correct identification of Monet paintings), along with a minimal number of misclassifications of only 8 false positives and 2 false negatives. This impressive performance is also measured in Figure 4, which lists excellent measures with regards to precision, recall, F1 score, and total accuracy. This performance highlights the model’s good performance in classical art verification tasks.

## 3. AI-Generated Art Detection Challenges

The model performed very poorly when evaluated on AI-generated paintings of Monet, as indicated in Figure 5. The confusion matrix indicates 118 false positives, in which AI-generated paintings were incorrectly classified as actual Monet works, compared to 18 that were classified correctly as non-Monet. Although the model performed well in the classification of actual works of Monet with 129 true positives, this was achieved with compromised overall classification accuracy. Figure 5 illustrates the difficulty of relying on metrics that display declining accuracy (0.5223), even when high recall performance is sustained (0.9485). The precipitate decline in overall accuracy to 0.5404, with a very low non-Monet accuracy of 0.1324, demonstrates the model’s issue in differentiating between AI-generated paintings and genuine paintings.



Total Accuracy	Precision	Recall	F1 Score	AI Accuracy
54.04%	52.23%	94.85%	67.36%	

Figure 5. Finalized results from the Monet Model comparing the Monet dataset vs the AI dataset, including the confusion matrix and performance metrics.



Figure 6. Visual comparison of the predictions made by the Monet model. The first row shows the model’s predicted probabilities for Monet (“M”) and non-Monet (“NM”) categories, with perfect accuracy. The second row presents non-Monet images, where the model assigns high probabilities to the non-Monet category, except for one case. The third row shows AI-generated images in Monet’s style, where the model’s accuracy is close to zero.



#### **4. Visual Analysis**

Figure 6 offers required insights through its visual juxtaposition of genuine Monet paintings, non-Monet paintings, and AI-generated paintings. The confidence scores (NM: Non-Monet, M: Monet) are the model's classification choices in various categories. Original Monet paintings all scored high confidence ratings (M: 0.96-1.00), and artificial intelligence-generated paintings also scored high confidence ratings (M: 0.92-1.00). This result indicates that AI-generated paintings have been successful in mimicking Monet's signature stylistic elements to an extent that poses challenge to existing classification techniques. Visual analysis especially brings out the advancement of contemporary AI art generation methods in imitating artist styles.

#### **5. Cross-Artist Validation**

The process was also validated on a Van Gogh dataset, with the same set of non-Monet paintings serving as a control group. The Van Gogh model scored 94% accuracy, aligned with the 96.32% performance of the Monet model on conventional artworks. Although the entire confusion matrix of the Van Gogh analysis was not saved due to technical constraints experienced by the study, the comparable accuracy metrics demonstrate the applicability of the binary classification method to the instance of traditional art authentication. Future research would be enhanced by a full replication of the Van Gogh analysis to generate comprehensive performance metrics in line with the ones presented for the Monet classification model.

#### **Conclusion**

This research indicates the potential and boundaries of

binary classification models in art authentication, in the new frontier of AI-generated artwork. This research yielded many significant findings that add to the field of computational art analysis and authentication.

The binary classification method has reported very high rates of efficiency in traditional art authentication, scoring 96.32% accuracy in Monet and 94% in Van Gogh attributed works. These findings confirm the feasibility of applying binary classification as an efficient means of scalable traditional art verification and indicate that one can create equivalent models for other artists with similar success rates.

This research also brought to light fundamental issues with AI-generated art. The model's evident struggle to distinguish between AI-generated artwork and authentic Monet paintings, with a 0.1324 accuracy rate for AI-generated art, is a fundamental challenge to the art authentication field. This finding highlights the sophisticated nature of modern AI art creation tools and suggests that traditional measures of appraisal might need to be modified in the case of AI-generated work.

The stark difference observed in the performance of the model in validating traditional art versus AI-generated art has strong implications for art authentication. With the increasing complexity level of AI-generated art, there is a vital need for authentication procedures to move beyond binary classification techniques. Future research needs to delve into multimodal authentication techniques that integrate other attributes with visual features, including provenance information, material analysis, and temporal consistency

The findings of this study also carry important implications for comprehending the nature of artistic style and authenticity in an era where artificial intelligence can simulate artistic

processes with increasing accuracy. The levels of confidence evident in artwork produced by AI imply that current computational techniques might need to be supplemented with new techniques that can capture the subtle aspects of artistic authenticity.

The current study suggests several productive avenues for future research. These include designing sophisticated neural network architectures specific to AI-generated art detection, combining different classification models within ensemble systems, and investigating unsupervised learning approaches to identifying characteristic artistic signatures that AI programs may struggle to replicate.

And that may not be enough. While the current research points towards directions such as developing sophisticated neural network architectures, ensemble classification models, and unsupervised approaches to detect artistic signatures, simply adding an “AI-generated” category to improve its classification is a feeble solution. All this is temporary relief to this ever-growing issue. With cutting-edge models such as DALL·E 3, Stable Diffusion, and Midjourney repeatedly pushing boundaries of what constitutes art, the line between human and machine creation is rapidly blurring. Instead of relying on a fixed binary framework, future research must embrace a more fluid, context-sensitive model and one that not only involves technical detection but also recasts our very concepts of authenticity as art itself becomes more liquid in the digital world.

Finally, binary classification has tremendous promise in traditional artwork authentication, the challenges presented by AI-generated works necessitate further development of both technological approaches and theoretical underpinnings to art authentication. As the intersection of art and artificial intelligence continues to develop, we

need to ensure our methods of artistic intent capture and authentication also move forward.

**Data Availability Statement:** The data and code for the training, validation, and test classes are available on GitHub via <https://github.com/bgorman3/ITSC-3990-AI-Art-Authentication/tree/main>

## References

1. Amineddoleh, L. (2018, May 10). How museums handle forgeries in their collections. Artsy. Retrieved from <https://www.artsy.net/article/artsy-editorial-museums-handle-forgeries-collections>
2. Dobbs, T., Nayeem, A.-A.-R., Cho, I., & Ras, Z. (2023). Contemporary art authentication with large-scale classification. *Big Data and Cognitive Computing*, 7(4), 162. <https://doi.org/10.3390/bdcc7040162>
3. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>
4. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, pp. 770-778, doi: 10.1109/CVPR.2016.90.
5. Hoving, T. (1996). *False Impressions: The Hunt for Big-Time Art Fakes*. Simon & Schuster.
6. Kavarakuntla, T., Han, L., Lloyd, H., Latham, A., & Akintoye, S. B. (2021). Performance analysis of distributed deep learning frameworks in a multi-GPU environment. 20th International Conference on Ubiquitous Computing and Communications (IUCC/CIT/DSCI/SmartCNS), 408-415. <https://doi.org/10.1109/IUCC-CIT-DSCI-SmartCNS55181.2021.00071>
7. Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
8. Ruiz, P. (2018, October 8). Understanding and visualizing ResNets. Towards Data Science. <https://medium.com/towards-data-science/understanding-and-visualizing-resnets-442284831be8>
9. Sabatelli, M., Kestemont, M., Daelemans, W., & Geurts, P. (2018). Deep transfer learning for art classification problems. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops* (pp. 0-0). [https://doi.org/10.1007/978-3-030-11021-5\\_48](https://doi.org/10.1007/978-3-030-11021-5_48)
10. Vincent, J. (2022, August 30). An AI-generated artwork won first place at a state fair fine arts competition, and artists are pissed. VICE. <https://www.vice.com/en/article/an-ai-generated-artwork-won-first-place-at-a-state-fair-fine-arts-competition-and-artists-are-pissed/>
11. Vincent, J. (2023, January 25). AI artist sues for copyright infringement, but the law isn't on their side. VICE. <https://www.vice.com/en/article/ai-artist-copyright-lawsuit/>
12. Dissection Layers. *Annals of Surgical Oncology*, 31(3), 1690–1691. <https://doi.org/10.1245/s10434-023-14633-7>
13. Rudzicz, F., & Saqur, R. (2020). Ethics of Artificial Intelligence in Surgery. arXiv.Org. <https://doi.org/10.48550/arxiv.2007.14302>
14. Shaik, T., Tao, X., Higgins, N., Li, L., Gururajan, R., Zhou, X., & Acharya, U. R. (2023). Remote patient monitoring using artificial intelligence: Current state, applications, and challenges. *Wiley Interdisciplinary Reviews. Data Mining and Knowledge Discovery*, 13(2), e1485-n/a. <https://doi.org/10.1002/widm.1485>
15. Siemionow, K., Katchko, K., Lewicki, P., & Luciano, C. (2020). Augmented reality and artificial intelligence-assisted surgical navigation: Technique and cadaveric feasibility study. *Journal of Craniovertebral Junction and Spine*, 11(2), 81–85. [https://doi.org/10.4103/jcvjs.JCVJS\\_48\\_20](https://doi.org/10.4103/jcvjs.JCVJS_48_20)
16. Zhou, X.-Y., Guo, Y., Shen, M., & Yang, G.-Z. (2020). Application of artificial intelligence in surgery. *Frontiers of Medicine*, 14(4), 417–430. <https://doi.org/10.1007/s11684-020-0770-0>

## Appendices

### Appendix A: Model Architecture Details

- Input layer: 224x224x3 RGB images
- Convolutional layers: Maintained original ResNet-101 structure
- Dropout rate: 0.3
- Final dense layer: 1 unit with sigmoid activation
- Total trainable parameters: 44.5M

### Appendix B: Training Parameters Final Hyperparameters used for Model Training

- Batch size: 32
- Total art: 2761 images
- Initial learning rate: 0.001
- Optimizer: Adam
- Loss function: Binary Cross-Entropy
- Early stopping patience: 5 epochs
- Learning rate reduction factor: 0.1
- Learning rate patience: 3 epochs
- Maximum epochs: 30
- Training/Validation/Test split: 70/20/10

### Appendix C: Dataset Composition

	Training	Validation	Test
Monet	957	273	136
Non Monet	977	279	139
AI Generated	0	0	136
Total	1934	552	411

### Appendix D: Detailed Training Data Breakdown

- <https://bit.ly/4ltOeMc>

### Appendix E: Technical Terminology

**Adam Optimizer:** An algorithm for optimization that modifies the learning rate in an adaptive manner for each individual parameter, enhancing its efficiency in training deep neural networks [7].

**Average Pooling:** Helps to reduce the size of feature maps by computing the average value of regions in an image, helping to retain important features while reducing computational complexity.

**Binary Classification:** A category of machine learning classifier whereby an algorithm acquires the ability to classify input data into two distinct categories or classes.

**Binary Cross-Entropy Loss:** Binary classification loss function that estimates the performance of a model with output as a probability value ranging from 0 to 1 [3].

**Confusion Matrix:** Matrix utilized to measure the performance of classification models by illustrating the number of correct and incorrect predictions by type.

**Deep Learning:** A subfield of machine learning which employs neural networks with many layers that are capable of learning representations of data at multiple levels automatically [3].

**Epochs:** Passing through the entire training dataset once while training a machine learning model.

**F1 Score:** A model accuracy measure that combines precision and recall into a single score.

**Generative Adversarial Networks (GANs):** This is a class of machine learning architectures where two adversarial neural networks try to produce new, synthetic instances of data that cannot be distinguished from real data [3].

**GPU Acceleration:** A process that uses a GPU to speed up computations by handling multiple tasks simultaneously, improving performance in deep learning and other intensive applications.

**Normalization:** Refers to the process of scaling values acquired on various scales to a single common scale, typically scaling data to a value between 0 and 1 [3].

**Overfitting:** is a model fault which happens when a machine learning model is overfit to training data, its noise, and variability, hence leading to poor performance on new data [3].

**Preprocessing Pipeline:** is a collection of data processing routines that are executed to raw data before its use in training a machine learning model [3].

**Provenance:** The origin, history and chain of custody of an object, document, or piece of art. It helps certify authenticity, ownership, and historical context by tracing where something comes from and how it is transferred overtime and any modifications that have been made or restorations that have been done.

**ResNet-101:** A deep neural network structure with 101 layers that is capable of training profound networks efficiently with residual connections [4].

**Tensor:** A multidimensional array which can store data in more than one dimension it is heavily used as the building data structure for deep learning frameworks [3].