Creativity and Innovation with Artificial General Intelligence: Designing Systems that Generate Novel Ideas

Nikhil Kassetty[1*], Srinivas Chippagiri[2], Karan Alang[3], Lakshman Kumar Jamili[4], Nuzhat Noor Islam Prova[5]

[1]Independent Researcher, Atlanta, GA, USA.
[2]Independent Researcher, Bellevue, WA, USA.
[3]Independent Researcher, Cupertino, CA, USA.
[4]Independent Researcher, Prosper, TX, USA.
[5]Independent Researcher, Queens, NY, USA.

*Corresponding author(s). E-mail(s): nikhilkassetty.cs@gmail.com; Contributing authors: cvas22@gmail.com; karan.alang@gmail.com; lakshman.techmind@gmail.com; nuzhatnsu@gmail.com;

Abstract

The combination of creativity and innovation is something that has always seemed to be stamped exclusively for humans — until recently, with the advances in Artificial General Intelligence (AGI). In this paper, a novel hybrid AI model involving combining Generative Adversarial Networks (GANs) and Efficient-NetB3 for increasing machine-driven creativity is presented. The model generates diverse artistic samples using GAN-based data augmentation, and can still perform accurate classification of artistic styles using EfficientNetB3. The approach was trained on the WikiArt dataset of 80,020 painted digitizations over 27 artistic styles. The performance of the hybrid model is experimentally shown to achieve a classification accuracy of 92.85% and exceeds the performance of existing AI-driven creative systems. The work presented in this research suggests that AI may be used to generate original artistic ideas, and it is joined by the larger discursive project surrounding AGI and creativity. Future work includes applying machine-generated creativity to other domains such as music, literature, and design, as well as assessing the novelty of machine-generated creativity.

Keywords: Creativity , Generative Adversarial Networks , EfficientNetB3 , Idea Generation , AI-Driven Innovation , Artificial Intelligence , AGI

R. Nagariya et al. (eds.), *Proceedings of the International Conference on Sustainable Business Practices and Innovative Models (ICSBPIM-2025)*, Advances in Economics, Business and Management Research 349,
https://doi.org/10.2991/978-94-6463-872-1_44

Human development has always depended on creativity and innovation, which has produced advances in the arts, sciences, and technology [1]. Subsequently, it has long been accepted that humans are capable of original thought, substantial artistic expression, and innovative concept development. However, because artificial intelligence (AI) is developing quickly, there is growing interest in creating systems that mimic and improve these creative processes [2]. The goal of AGI, which aims to carry out intellectual tasks comparable to human cognition, has become a significant area of investigation [3]. Despite advancements in deep learning (DL) [4], developing AI systems that innovate rather than mimic existing patterns remains a considerable challenge [3]. Many current techniques struggle to move beyond data-driven replication, failing to produce novel and diverse creative outputs [5].

Existing AI models for creativity often face limitations in generating unique content [6]. Conventional GANs and Variational Autoencoders (VAEs) usually produce outputs that closely replicate training data instead of inventing new innovative techniques despite their effectiveness in creating realistic illustrations [7]. This lack of creativity results from overfitting to existing patterns instead of pursuing the entire range of creative alternatives. However, the classification methods utilized to analyze novel concepts tend to decline, identifying subtle stylistic changes and leading to a lack of cross-genre generalization. The availability of diverse and high-quality data sets that accurately reflect cultural influences and creative genres is another significant barrier to the creativity generated by AI. The scarcity of such datasets restricts AI models' proficiency in comprehending and producing innovative ideas.

To overcome these drawbacks, this study presents a hybrid strategy combining EfficientNetB3 for classification with GANs for data augmentation. GANs generate synthetic artworks, expanding the diversity of the dataset and lowering duplication in training samples. We employed the WikiArt dataset for this study due to its expansive collection of paintings from different genres, historical periods, and creative ideas. This study presents a robust and scalable approach to designing intelligent systems capable of producing and recognizing novel artistic ideas by addressing the fundamental challenges of AI-generated creativity. By integrating generative and classification techniques, this study progresses the field of AI-driven creativity and lays the groundwork for future explorations into machine-assisted innovation.

The main contribution of this study is as follows:

1. We present a novel hybrid GAN+EfficientNetB3 model that effectively overwhelms existing limitations in AI-generated creativity. Unlike traditional models that struggle with generalization and diversity, our proposed architecture synergizes the strengths of generative and classification-based DL techniques, resulting in a more robust creative AI system.
2. We employed the WikiArt dataset with 25 classes and a collection of 80,020 unique digitized paintings created by 1,119 artists across 27 distinct artistic styles. We also improve dataset variability through GAN-based augmentation, providing the generated artistic outputs indicate higher originality and decreased redundancy

3. Our model enhances classification accuracy, differentiating artistic styles and novel creations better. It also contributes to the broader discourse on AGI and creativity by demonstrating a scalable, adaptable framework that bridges the gap between machine learning (ML)-based artistry and human-like innovation.

The rest of this paper is structured as follows: In section 2, the applications of several AGI techniques and novel ideas are reviewed in the literature. We describe the methodology in Figure 1, including dataset selection, preprocessing techniques, and our proposed model architectures. The experimental results and performance analysis are presented in section 4. Finally, insights and future research directions are concluded in section 5.

2 Literature Review

Life has always been enriched with creativity and innovation, and now AGI is being explored for its role in provoking new ideas by researchers. AI-powered generative models complement some traditional methods of creative problem solving like brainstorming and design thinking through the help of AI-powered generative models for design synthesis and the help of generative models in concept development. It is through recent studies that the potential of AI to enhance human creativity has been demonstrated in different domains such as those pertaining to fashion designing, product development, and co-creative design systems. In this literature review, we explored bare highlights of past developments in AI-driven creativity, specially taking into account how generative models and AI systems are evolving consciousness ideation procedures and helping human designers create innovative answers.

Sbai et al. [8] employed GANs to generate original and creative fashion designs and introduced a creativity loss function to add novelty to the generated fashion designs. Building on automated metrics and human studies, they evaluated them, which indicated that 61% of these AI-generated designs were considered to be of human nature. These results show that AI can assist human creativity in fashion design because their proposed loss outperformed Creative Adversarial Networks (CANs), and in novelty and likability, the highest scores were achieved. Karimi et al. [9] offered a computational model of conceptual shifts using novelty on DL-based vector representations. On the experimental side, they used CNN-LSTM architectures, in which each sketch was represented as the final LSTM layer of 256 values per sketch. There was effective learning and the accuracy plateaued at 73.4% after 1 million training steps. This finding suggests that ML is enhancing human artistic innovation by augmenting the creative process through the use of AI-driven design tools that assist, but also very much add to the creative process.

Nobari, Rashad, and Ahmed [10] proposed a novel generative network that addresses the limitations of existing GAN in creative design synthesis through a new concept and forms a new kind of generative network. They presented interesting anecdotes about deep neural networks' role in several fields, such as image recognition, and language understanding, where they seek to identify patterns in high-dimensionality datasets and perform predictions with high accuracies. The study recognizes the possibility of generative models, in particular GANs, to enable human designers to facilitate the exploration of the iterative design process. Kim, Maher, and Siddiqui [11] also explored how AI-based co-creative design systems can amplify designers' creative process to collaborate with an AI agent. The study centers around open-ended creative tasks with an artificial intelligence-driven approach to guide the initial stages of idea generation by offering inspirational design solutions. They propose a system that generates those solutions based on visual and conceptual closeness to sketches drawn by designers, providing an interactive and dynamic creative experience. Vocke, Constantinescu, and Popescu [12] explored how digital technology provides a transformative impact on economies and societies taking advantage of increasing artificial intelligence (AI), robotics, big data, business analytics, and extended reality. This study stresses the importance of these technologies that help total digital connectivity and better data analysis that will improve their productivity and result in optimized business performance. While integrating AI into the design and development process, as the authors purport it, validates the importance of human creativity and harnesses the power of machine intelligence for efficiency and innovation, this survey is an attempt to present available-based technological solutions that aid inhuman-machine collaboration in creative and design based activities. Rick et al. [13] showed how generative AI, such as GPT-3.5, can be used to generate ideas using the system, Supermind Ideator, a system that leverages generative AI to help idea generation. Some of the work they build on is traditional creativity support methods (suggested as structured approaches such as brainstorming and design thinking and as software tools for facilitating the recording and sharing of ideas). Now, generative AI can provide such a transformative shift not only to helping document, the authors say, but actually to thinking about the work itself and offering up novel ideas to the user they might not have thought of on their own.

3 Methodology

In this section, we will briefly discuss our dataset and data preprocessing techniques that we employed along with data augmentation. After that, we discuss our methodology in-depth and present their strengths. Figure 1 depicts the overall research methodology.

3.1 Dataset Description

For this study, we utilized the WikiArt dataset [14], a publicly available collection of 80,020 unique digitized paintings created by 1,119 artists across 27 distinct artistic styles, including Romanticism, Baroque, and Impressionism, and 10 painting genres, such as landscape, portrait, and still life which includes artwork from the 15th to 20th centuries, offers a wealth of information for stylistic analysis, artistic development, and innovative AI applications. The structured metadata is very useful for training AI models in style classification and generative tasks since it allows for accurate categorization. The pictures from WikiArt.org make it possible to comprehend artistic variation and creative patterns on a deeper level. It is essential to acknowledge that WikiArt is available exclusively for non-commercial research, with usage governed by the terms and conditions of WikiArt.org. Figure 2 presents the dataset sample, and

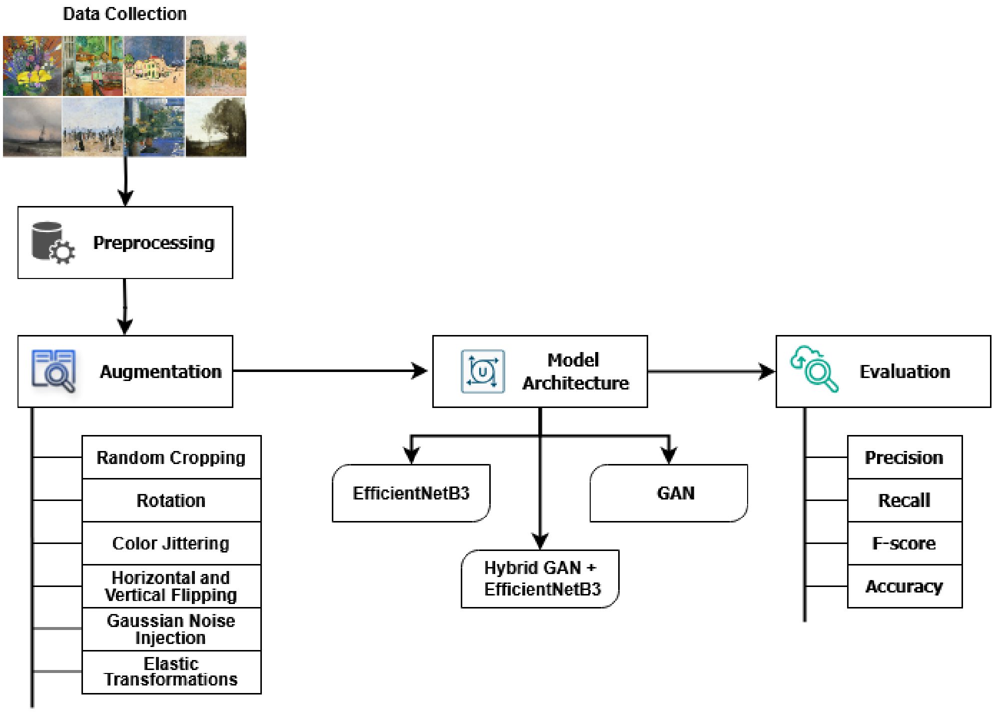

Fig. 1: Graphical Representation of our Overall Research Methodology

Figure 3 displays the number of images per style with varying counts across different categories.

Fig. 2: Sample of the dataset

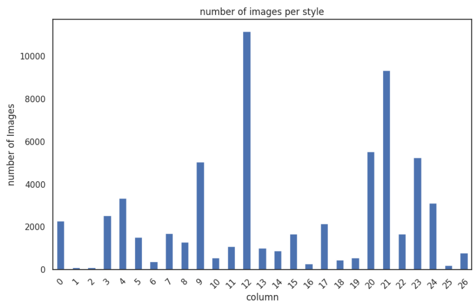

Fig. 3: Number of images per style, with varying counts across different categories

3.2 Dataset Preprocessing and Standardization

Preprocessing the data is an essential step in implementing a robust model. The training and validation datasets were initially looked at for missing values, and their incorrect filenames were removed to provide consistency. The size of each image was formatted to 256 x 256 x 3 for consistent format throughout the dataset. We applied a concatenation-based expansion method for images smaller than this resolution, duplicating and merging the image before center-cropping it to the required size. To further refine the dataset, near-duplicate images—often the result of multiple digitizations of the same artwork—were detected using perceptual hashing and subsequently removed.

3.2.1 Balanced Sampling Strategy:

Working with historical art collections has several difficulties, including the unequal representation of artists and genres; some have a disproportionately high number of pieces. This was addressed using a weighted sampling technique, and the overrepresentation of dominant artists was avoided. Our final data set maintained a balanced representation across 27 artistic styles, with approximately 6,000 images per category, enabling diversity and lowering bias in model training. After balancing, we prepared the dataset into 50,000 training and 10,000 testing images, providing a well-distributed foundation for model evaluation.

3.2.2 Standardized Formatting:

We formatted the dataset into considerable resolutions for compatibility with different ML frameworks: 32×32, 256×256, and original image sizes. The 32×32 and 256×256 versions were developed using a center-cropping technique. We minimized artifacts while downsampling using the Nicolas Robidoux resampling technique, which preserved structural coverage and fine details. The final dataset was organized into class-specific folders using the ImageFolder format (256 x 256) to facilitate easy incorporation into DL workflows.

3.3 Data Augmentation

Augmentation was essential to enhancing generality and avoiding the model overfitting to particular patterns because the WikiArt dataset includes historical artworks from various periods and styles. We synthesized extra training examples using an augmentation technique based on GANs to capture complex artistic qualities across many painting styles. Furthermore, we combined traditional augmentation methods with GAN-based augmentation, including:

1. Random Cropping: Removing distinct portions of an artwork to enrich spatial variance.
2. Rotation: Improving invariance to exposure transformations.
3. Color Jittering: Brightness, contrast, and saturation can be modified randomly to adjust differences in lighting and artistic palettes.
4. Horizontal and Vertical Flipping: Ensuring the model recognizes artistic elements regardless of orientation.

5. Gaussian Noise Injection: Simulating imperfections found in historical digitized artworks to improve resilience against noise.
6. Elastic Transformations: Applying minor distortions to mimic variations in brush strokes and artistic textures.

These augmentation techniques expanded the dataset, doubling the training samples and maintaining visual and stylistic consistency.

3.4 Model Architecture

3.4.1 EfficientNetB3

This study employs EfficientNetB3 as the primary feature extraction and classification model due to its efficient scaling strategy and superior performance in high-resolution images [15]. Using a compound scaling approach, EfficientNetB3 modifies the number of layers, number of filters per layer, and input resolution of the network to optimize computing efficiency without compromising accuracy [16]. The architecture incorporates stacked convolutional layers, observed by batch normalization and ReLU activation functions, to stabilize training and improve convergence. The model utilizes depthwise separable convolutions, which perform channel-wise operations and substantially reduce the number of parameters while maintaining representational power.

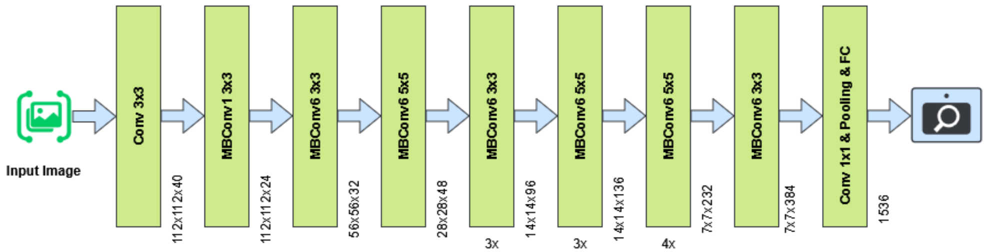

Fig. 4: EfficientNetB3 Architecture

We employed EfficientNetB3's pre-trained weights for feature extraction, which have been refined on extensive image datasets. The model's extraction of multi-scale hierarchical features captures high-level semantic information and low-level textures crucial for categorization. The last layers are a fully linked layer for classification and a global average pooling layer that decreases the size of the feature map. This architecture ensures faster inference times, reduced memory consumption, and improved generalization across our dataset's diverse artistic styles. Figure 4 graphically represents the EfficientNetB3 architecture.

3.5 Generative Adversarial Network (GAN)

To increase the generalization and durability of the model, we supplemented the training dataset with synthetic GAN-generated samples. Realistic samples are produced by the Generator G, and generated samples are separated from actual data by the Discriminator D. The goals of the adversarial training procedure are as follows:

$$\min_G \max_D \mathbb{E}_{x \sim P_{\text{data}}}[\log D(x)] + \mathbb{E}_{z \sim P_z}[\log(1 - D(G(z)))] \quad (1)$$

where, P_{data} defines the distribution of real data, and P_z is the previous distribution of the latent variable.

We applied filtering strategies based on the Structural Similarity Index (SSIM) [17] and Fréchet Inception Distance (FID) [18] to ensure high-quality synthetic samples. The SSIM score is calculated as follows:

$$SSIM(x, y) = \frac{(2\mu_x \mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (2)$$

where, μ_x and μ_y are mean intensity values, σ_x^2 and σ_y^2 represent variances, and σ_{xy} is the covariance between images x and y. Only images with $SSIM > 0.9$ were retained.

Filtered GAN-generated images were mixed with actual samples using a weighted sampling approach. The final loss function integrated a regularization term for synthetic data:

$$L = L_{\text{CE}} + \lambda L_{\text{GAN}} \quad (3)$$

where, L_{CE} is the cross-entropy loss for classification, and L_{GAN} regulates the effect of generated data. This integration process improved the model's ability to handle diverse image distributions while maintaining classification accuracy.

3.6 Proposed Hybrid GAN + EfficientNetB3 Model

In this study, we introduce a novel hybrid architecture that combines the robust feature extraction and classification capabilities of EfficientNetB3 with the generative capabilities of GAN. The model incorporates EfficientNetB3 and a GAN to identify generated and authentic images to produce synthetic images. The hybrid method teaches the GAN's generator to create artificial images from the dataset that mimic actual photos. Combined with the actual images, these images are sent into EfficientNetB3, which is tuned for feature extraction and classification. The generator G of the GAN is responsible for creating synthetic data $\hat{x}$. At the same time, the discriminator D differentiates between original images x from the dataset and the synthetic images generated by G. The generator and discriminator are trained simultaneously, optimizing the adversarial loss function L_{GAN}, which facilitates the generator to produce realistic images. The adversarial loss function is expressed as:

$$L_{\text{GAN}} = -\mathbb{E}_{x \sim p_{\text{data}}(x)}[\log D(x)] - \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))] \quad (4)$$

where, the distribution of actual data is represented by $p_{\text{data}}(x)$ and the distribution of random noise supplied into the generator by $p_z(z)$, the discriminator D improves its ability to discern between genuine and false pictures, while the generator G learns to produce synthetic images that are identical to actual photos.

Once the synthetic images are generated, they are combined with the actual dataset, forming an augmented training set. The EfficientNetB3 model, which is intended to extract features from the photos and carry out classification, is trained using this enhanced database. Depth-wise separable convolutions are used by EfficientNetB3 to effectively extract features while reducing the model size and increasing computing performance. EfficientNetB3's categorization process may be shown as follows:

$$f_{\text{EffNetB3}}(x) = \text{Softmax}\left(W \cdot \text{GlobalAvgPool}\left(\text{Conv}(x)\right) + b\right) \tag{5}$$

where, x is the input image, $\text{Conv}(x)$ represents the sequence of convolution operations used to x, and GlobalAvgPool represents the global average pooling operation. The weights W and biases b in the fully connected layer produce the final class probabilities.

To optimize the hybrid model, we define a combined loss function that includes both the adversarial loss from the GAN and the classification loss from EfficientNetB3. The total loss L_{total} is a weighted sum of the cross-entropy loss L_{CE} for classification and the adversarial loss L_{GAN}:

$$L_{\text{total}} = \lambda_1 L_{\text{CE}} + \lambda_2 L_{\text{GAN}} \tag{6}$$

where, $L_{\text{CE}} = -\sum_{c=1}^{C} y_c \log(p_c)$ is the cross-entropy loss, with y_c being the true label and p_c being the predicted class probability for class c, L_{GAN} is the adversarial loss from the GAN, and λ_1 and λ_2 are hyperparameters controlling the classification's relative importance and adversarial losses. The generator G and EfficientNetB3 are jointly optimized by backpropagating the gradients through both networks. This simultaneous optimization ensures that the generator learns to produce high-quality synthetic images while EfficientNetB3 effectively classifies real and synthetic data.

During training, the GAN is first trained to generate realistic images by minimizing the adversarial loss. Once the GAN has produced sufficiently realistic synthetic images, the augmented dataset is fed into EfficientNetB3. EfficientNetB3 learns to classify both real and synthetic images, with the classification loss guiding the model to output accurate predictions. The tot l loss function L_{total} is minimized using an optimization algorithm such as Adam, which adapts the learning rates for each parameter during training.

4 Results & Discussion

4.1 GAN Training Performance Analysis

To improve the diversity and creativity of the dataset, a data augmentation was performed with the help of a GAN. The GAN has a Generator that generates synthesized images and a Discriminator that labels if a given image is real or synthetic. Here, the network of the Generator trains to produce realistic samples against the network of the Discriminator to discriminate real and synthetic images.

Figure 5 shows Generator and Discriminator's loss curves over the training epochs. The Discriminator starts out with a very high loss value and it is unable to differentiate between real and synthetic images. Nevertheless, the loss for the Discriminator becomes stable whereas the Generator's loss is growing in the meantime, reflecting its

Fig. 5: Discriminator and Generator loss curves during GAN training.

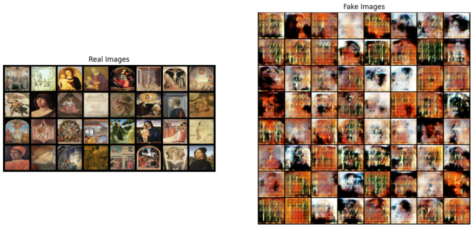

Fig. 6: Comparison between original dataset images and GAN-augmented images.

better generation of real-looking samples. Both losses involve an oscillatory nature, and they depend on the adversarial learning process along which both networks are adaptively modified according to each other's progress.

A qualitative comparison between the original dataset images and the GAN-augmented images is conducted to evaluate the quality of the generated images. In Figure 6, some of the real and synthetic images are presented. A variation of artistic features is augmented in the images, creating more diversity in the training data, while preserving the stylistic elements of the original artworks.

4.2 Hybrid Model Performance Analysis

The training and validation accuracy curves of our proposed hybrid GAN + EfficientNetB3 model are shown in Figure 7. The red solid line indicates that the training accuracy increases steadily until the near-perfect classification performance. In particular, the validation accuracy (blue dashed line) has steadied ups and downs in the early epochs; and finally attains a high accuracy level. The ultimate testing accuracy of our proposed model constitutes 92.85%, revealing its accuracy in classifying artistic styles from the WikiArt dataset. At the beginning of the epochs, the validation accuracy oscillates because it's an artistic image dataset, and it's impossible for the model to generalize across different styles.

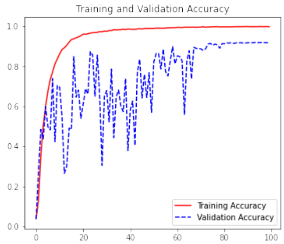

Fig. 7: Training and validation accuracies of the proposed hybrid model over epochs.

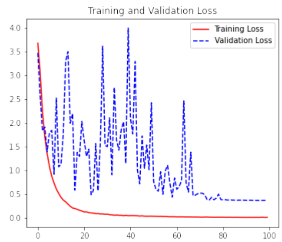

Fig. 8: Training and validation losses of the proposed hybrid model over epochs.

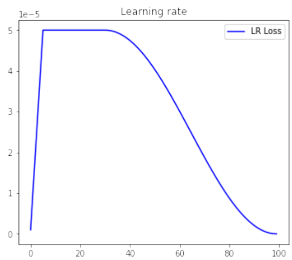

Fig. 9: Learning rate schedule during training.

Further validations of the effectiveness of our proposed model are achieved with the loss curves shown in Figure 8. The training loss (red solid line) also drops gradually, so we can say that the model does learn effectively. Meanwhile, the validation loss (blue dashed line) has severe fluctuations in the initial training phase, tapering down with a gradual increase in training. The pattern here implies that the model is not strong with such a diverse dataset, but the model eventually stabilizes and shows a strong ability of generalization of all kinds of art.

The learning rate schedule used in the training process is illustrated in Figure 9. The curve in its form is similar to a typical cyclical pattern. With time, the learning rate does initially go up to a peak and then slowly drops down. This adaptive learning rate technique guarantees to update the weights adaptively such that the weight updates are optimal, preventing premature convergence and avoiding overfitting risk at the same time. This smooth decrease of the learning rate to the later epochs aids the model's performance to arrive at a high classification accuracy.

4.3 Precision Analysis

This metric precision is critical when passing on the classification model's ability to perform because it relates the portion of the predicted sample of one given class that was properly classified against the entirety of it. The score means that there is a higher precision, that is, the model is very confident in its prediction and did fewer false positives. The precision of the proposed model across most artist classes is consistently high and always between 0.75 and 0.98, which is shown in Figure 10.

The precision is highest (0.98) for 'Gustave Doré' where the model is able to identify his works with very few false positives. Much like 'Eugène Boudin', 'Ivan Aivazovsky', and 'Rembrandt', precision scores over 0.90 are achieved, indicating their individual artistic characteristic are easy to recognize. However, the precision of 'Ilya Repin' (0.75) is the lowest, as the radicals and his illustrations are quite similar to the works of other realist artists. Minor variations were found in the results, however, the overall precision scores corroborate the fact that the hybrid model is capable of distinguishing artworks with high confidence.

4.4 Recall Analysis

The ability of the classification model to identify all relevant instances of a particular class is known as recall. A high recall score implies that the model obtains a good combination of the true samples with the least number of false negatives. As shown in Figure 11, recall scores for different artist classes ranges from 0.68 to 0.97. The model manages to retrieve, to the highest recall values (0.97) all artworks by 'Gustave Doré' and 'Eugène Boudin', i.e., it has almost all artworks on hand. Secondly, the model recalls 'Nicholas Roerich' and 'Ivan Aivazovsky' over 0.94, demonstrating that the model is capable of recognizing these artistic styles.

However, the recall for 'Salvador Dalí' is the lowest recall score (0.68) points out that some of his work may be misidentified due to their resemblance to surrealist or abstract works by other artists. The model has minor fluctuation in certain classes,

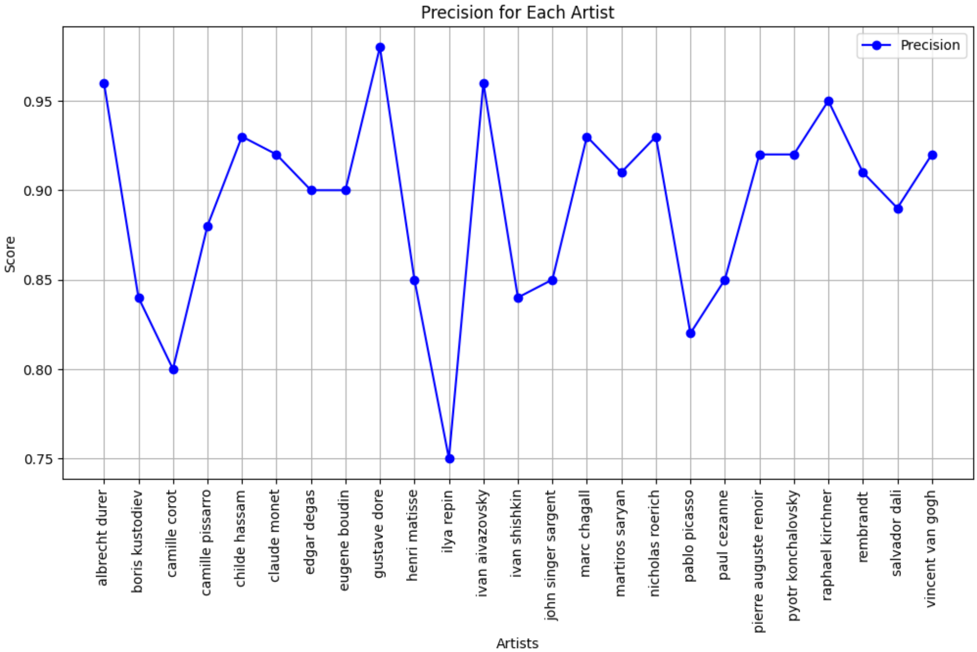

Fig. 10: Line plot showing precision scores for each artist class.

but it achieves strong retrieval performance for most artists, and hence we have a high proportion of the correctly retrieved samples.

4.5 F1-score Analysis

Precision and recall are balanced by the F1-score and thus combined into a measure of a model's performance. This is especially useful when dealing with class imbalances so that neither false positives nor false negatives will dominate the evaluation. As depicted in Figure 12, the F1-score lies in [0.74, 0.98], which shows that the model is not struggling between precision and recall. 'Gustave Doré' reaches the highest F1-score of (0.98), then followed by 'Eugène Boudin' (0.94), 'Ivan Aivazovsky' (0.96), and 'Nicholas Roerich' (0.94).

In other words, these scores mean that the model is able to classify these artists well enough to have minimal errors for this task. On the contrary, 'Ilya Repin' and 'Salvador Dalí' achieve the lowest F1-scores, 0.74 and 0.77 respectively, as they most likely share stylistic similarities with other artists in the dataset. With these variations, however, the high F1-scores still indicate that the hybrid model accurately classifies a broad category of art styles.

4.6 Comparative Analysis

?? compares our proposed GAN+EfficientNetB3 hybrid model with existing approaches for artistic style classification. In the case of the Painting 91 dataset, Chu

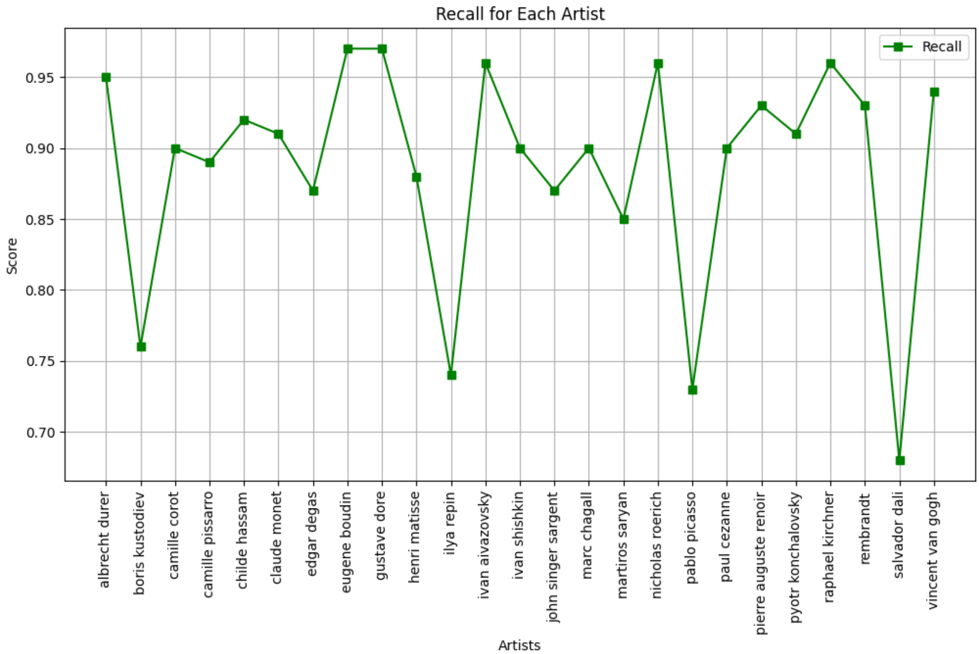

Fig. 11: Line plot showing recall scores for each artist class.

and Wu [19] applied a CNN with an accuracy of 64.32% on 91 classes. In their work [20], Bianco et al. used a DNN on MultitaskPainting100k that consisted of 1508 classes, but they had only gotten as much as 56.5% accuracy. Centic, Lipic and Grgic [21] applied a CNN to the WikiArt dataset containing 23 classes with 81.94% accuracy.

On the other hand, our proposed hybrid GAN+EfficientNetB3 model trained on WikiArt (25 classes) gets a considerably higher accuracy of 92.85%. This amounts to a 10.91% improvement from the Cetinic, Lipic, and Grgic's model, 28.53% from Chu and Wu, and 36.3% from Bianco et al. These superior performances demonstrate how the augmentation based on GAN can be effectively used, and how efficient feature extraction of the EfficientNetB3 can better generalize the artistic styles.

5 Conclusion

In this study, a novel hybrid AI model of GANs and EfficientNetB3 was introduced to improve machine-generated creativity. This approach addressed the weak points of conventional AI models, which tend to reproduce training data rather than truly creative content when they are trained and classified, and when they try to generate artistic styles. The model performs well on the WikiArt dataset with a maximum classification accuracy of 92.85% in the area compared to the previous AI-based artistic classification method.

Therein, the research illustrates the learning opportunity that comes from combining generative and classification techniques to set the bar further for AI-assisted

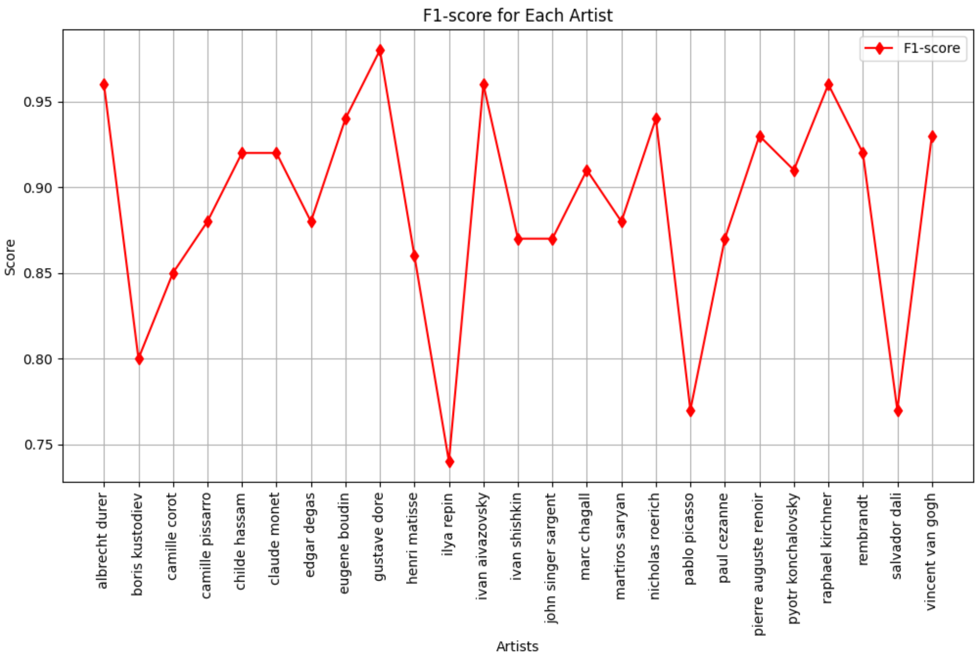

Fig. 12: Line plot showing F1-scores for each artist class.

innovation. Dataset diversity was contributed by the GAN component with their synthetic but diverse artistic works, and the EfficientNetB3 was used to perform style classification with maximum precision. It further showed difficulties in separating artists with very similar styles and good precision and recall scores in most categories.

The aim of future research could aim at integrating reinforcement learning in the creativity assessment for further enhancement and a wider application of the model to other creative domains like music or literature. Moreover, there is an open challenge towards refining the AI-driven evaluation metrics for novelty as well as originality in the domain of creative AI and AGI.

References

[1] P. D. R. Bambi, J. B. B. Pea-Assounga, Unraveling the interplay of research investment, educational attainment, human capital development, and economic advancement in technological innovation: A panel var approach, Education and Information Technologies 30 (3) (2025) 3309–3341.

[2] M. E. Yawar, M. Q. Hakimi, The impact of artificial intelligence technology on human resources performance in organizations, EuroGlobal Journal of Linguistics and Language Education 2 (1) (2025) 96–108.

[3] R. Raman, R. Kowalski, K. Achuthan, A. Iyer, P. Nedungadi, Navigating artificial general intelligence development: societal, technological, ethical, and brain-inspired pathways, Scientific Reports 15 (1) (2025) 1–22.

[4] S. S. Khan, A. U. H. Rupak, W. W. Faieaz, S. Jannat, N. N. I. Prova, A. K. Gupta, Advances in medical imaging: Deep learning strategies for pneumonia identification in chest x-rays, in: 2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT), IEEE, 2024, pp. 1–7.

[5] N. Henke, L. Jacques Bughin, The age of analytics: Competing in a data-driven world (2016).

[6] M. B. Garcia, The paradox of artificial creativity: Challenges and opportunities of generative ai artistry, Creativity Research Journal (2024) 1–14.

[7] E. K. Innocent, Enhancing data security in healthcare with synthetic data generation: An autoencoder and variational autoencoder approach, Master's thesis, Oslo Metropolitan University (2024).

[8] O. Sbai, M. Elhoseiny, A. Bordes, Y. LeCun, C. Couprie, Design: Design inspiration from generative networks, in: Proceedings of the European conference on computer vision (ECCV) workshops, 2018, pp. 0–0.

[9] P. Karimi, M. L. Maher, N. Davis, K. Grace, Deep learning in a computational model for conceptual shifts in a co-creative design system, arXiv preprint arXiv:1906.10188 (2019).

[10] N. N. I. Prova, Healthcare fraud detection using machine learning, in: 2024 Second International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI), IEEE, 2024, pp. 1119–1123.

[11] J. Kim, M. L. Maher, S. Siddiqui, Collaborative ideation partner: Design ideation in human-ai co-creativity., in: CHIRA, 2021, pp. 123–130.

[12] N. N. I. Prova, Enhancing fish disease classification in bangladeshi aquaculture through transfer learning, and lime interpretability techniques, in: 2024 4th International Conference on Sustainable Expert Systems (ICSES), IEEE, 2024, pp. 1157–1163.

[13] N. N. I. Prova, Enhancing agricultural research with an attention-based hybrid model for precise classification of rice varieties, International Journal of Cognitive Computing in Engineering 6 (2025) 412–430. doi:10.1016/j.ijcce.2025.02.002.

[14] T. Yemelianenko, I. Tkachenko, T. Masclef, M. Scuturici, S. Miguet, An approach for dataset extension for object detection in artworks using open-vocabulary models, in: Proceedings of the European Conference on Computer Vision (ECCV) Workshops, 2024.

[15] J. Spoorthi, K. Keerthi, S. Divya Sankari, V. Pandiyaraji, et al., Optimizing bone fracture diagnosis: A hybrid approach using efficientnet-b3 and svm for accurate classification, in: 2024 International Conference on Emerging Research in Computational Science (ICERCS), IEEE, 2024, pp. 1–7.

[16] A. R. Gadah, S. Isaak, A.-M. H. Saad, N. A. Mokhtar, Revolutionizing plant health: Leveraging efficientnet b3 for advanced tomato disease detection in agricultural ai, in: 2024 IEEE Symposium on Industrial Electronics & Applications (ISIEA), IEEE, 2024, pp. 1–6.

[17] I. Bakurov, M. Buzzelli, R. Schettini, M. Castelli, L. Vanneschi, Structural similarity index (ssim) revisited: A data-driven approach, Expert Systems with

Applications 189 (2022) 116087.
[18] Y. Yu, W. Zhang, Y. Deng, Frechet inception distance (fid) for evaluating gans, China University of Mining Technology Beijing Graduate School 3 (11) (2021).
[19] W.-T. Chu, Y.-L. Wu, Image style classification based on learnt deep correlation features, IEEE Transactions on Multimedia 20 (9) (2018) 2491–2502.
[20] S. Bianco, D. Mazzini, P. Napoletano, R. Schettini, Multitask painting categorization by deep multibranch neural network, Expert Systems with Applications 135 (2019) 90–101.
[21] E. Cetinic, T. Lipic, S. Grgic, Fine-tuning convolutional neural networks for fine art classification, Expert Systems with Applications 114 (2018) 107–118.