A Deep Learning Approach Based on Image Patch Sets for Art Forgery Detection

Soonchul Jung, Jae Woo Kim, Jin-Seo Kim, Yoon-Seok Choi†
Content Research Division
Electronics and Telecommunications Research Institute
Daejeon, Korea
{s.jung, jae_kim, kjseo, †ys-choi}@etri.re.kr

***Abstract*—With the recent expansion of the art auction market, the discernment of forged artworks has become increasingly vital. Studies have been conducted to detect forgeries through various means, such as physical examinations of paint and canvas, as well as more abstract inquiries into the artistic style. Among these, style-based studies have faced challenges due to the lack of relevant datasets. To address this, we have constructed a dataset by manually creating both genuine and forged oil paintings. Typically, artwork images are very large. Previous research has extracted small image patches for input but often failed to represent the artwork's features, depending on the patch's location within the work. In this paper, we propose a deep learning approach that utilizes a set of patches instead of a single image patch to determine whether the given artworks are from the same artist. Using multiple image patches is advantageous because they can represent the characteristics of the artwork more effectively than a single image patch. Experimental results demonstrate that the proposed approach achieves an accuracy ranging from 76% to 99%, with the accuracy increasing as the image patch set size grows.**

***Index Terms*—Art Forgery Detection, Pretrained ResNet, Feature Extractor, Image Patch Set, Forgery Dataset**

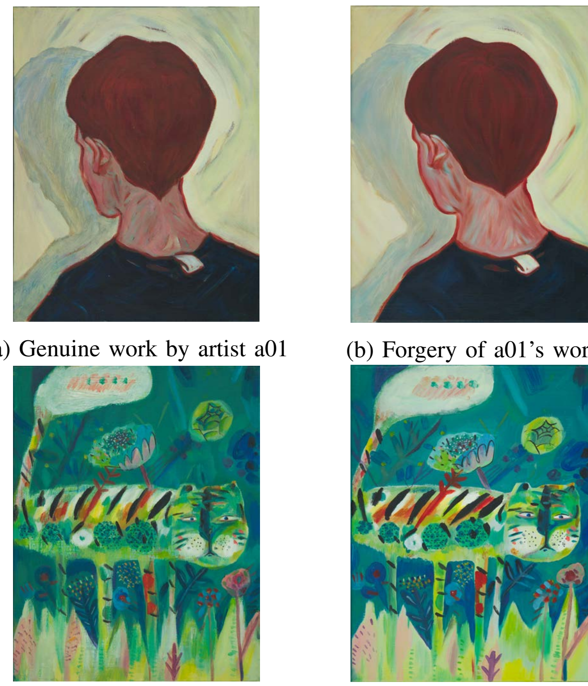

(a) Genuine work by artist a01 (b) Forgery of a01's work

(c) Genuine work by artist a02 (d) Forgery of a02's work

Fig. 1: Comparison of genuine works and forgeries

I. Introduction

The art world has witnessed a significant expansion in the auction market in recent years. This growth has not only fueled interest in art collection but has also led to an increase in the prevalence of art forgeries. The ability to accurately discern genuine artworks from forgeries has thus become a critical concern for artists, collectors, galleries, and auction houses alike.

Historically, the detection of art forgery has been a complex and multifaceted task. Traditional methods have ranged from physical examinations of paint and canvas to more abstract inquiries into the artistic style. While physical examinations provide concrete evidence [1], they can be invasive and potentially damaging to the artwork. On the other hand, investigations into artistic style, though less intrusive, have faced challenges due to the lack of relevant datasets.

Recent advancements in technology have opened new avenues for art forgery detection and artist identification. Noord *et al.* proposed a CNN called PigeoNET that automatically recognizes artists by their artworks [2]. Utilizing a large collection of digitized artworks, PigeoNET achieves an accuracy of over 70% in attributing previously unseen artworks to the correct artists. The network also provides insights into artist-specific characteristics within spatial regions of the artworks, representing a promising approach for computer-supported examination of art. Hwang *et al.* authors propose an optical analysis system to protect paintings from counterfeiting [3]. Elgammal *et al.* introduced a computational method to quantify individual stroke characteristics in line drawings [4]. This approach encompasses both global and local shape features, along with a deep neural network to capture variations in the local shape and tone of each stroke. The study involved a comparison of various feature types and presented results at both the stroke classification and drawing classification levels. Chen *et al.* utilized Convolutional Neural Networks (CNNs) in the authentication of works by Portuguese artist Amadeo de Souza Cardoso [5]. The study revealed that neural networks significantly outperformed traditional algorithms, even with limited samples. Ji *et al.* extended machine learning analysis to the surface topography of painted works [6]. A controlled

Fig. 2: A set of 80 images capturing artist a02's artwork

study was designed with paintings produced by art students, and the paintings were scanned using a confocal optical profilometer to produce surface height data. The surface data were divided into virtual patches and used to train CNNs for attribution. The resulting attribution was found to be 60 to 96% accurate. Kim *et al.* analyzed the 3D morphology of cracks in oil paintings to distinguish forgeries from original artworks, revealing distinct differences between original and fake cracks [7].

Artworks are often photographed at ultra-high resolutions to capture the finest details. However, conventional neural networks like CNNs typically handle images of around 256x256 pixels. To fit this size constraint, researchers usually take random crops from the high-resolution artwork images, creating very small patches for analysis. But there's a challenge: if these patches are taken from a simple background or are too small relative to the overall size of the artwork, they might not contain the unique features needed to distinguish one piece from another. Previous studies have often relied on single image patches [2], [5], [6], limiting the ability to accurately represent and authenticate artworks.

In response to these challenges, we propose a novel deep learning approach that utilizes multiple image patches instead of a single one. By leveraging a set of patches, we aim to represent the characteristics of the artwork more effectively and determine whether the given artworks were created by the same artist. This approach is grounded in the belief that a larger set of image patches can provide a more comprehensive view of the artwork, thereby enhancing the accuracy of forgery detection.

II. Dataset Creation

Our research necessitates a robust dataset comprising both authentic paintings and their corresponding forgeries for the purpose of discerning genuine works from imitations using deep learning algorithms. However, the challenge lay in the scarcity of an existing dataset that sufficiently met these specific criteria.

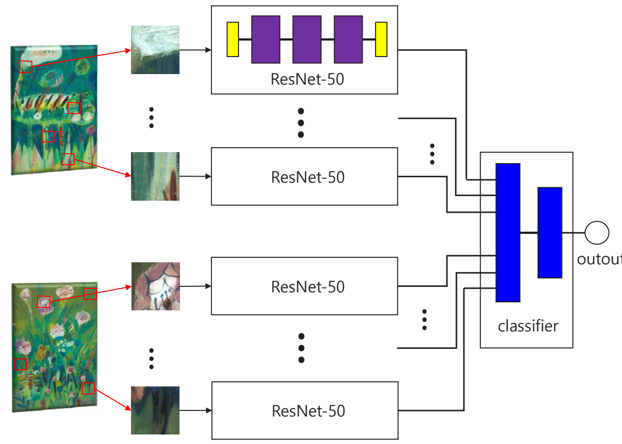

Fig. 3: Proposed Network Structure

To overcome this obstacle, we embarked on the creation of a unique dataset. We commissioned five professional artists with over 10 years of experience, referred to as a01, a02, a03, a04 and a05, to produce six paintings each, reflecting their individual styles. These paintings served as the authentic works, or "genuine paintings," and amounted to a total of 30 pieces. All the artworks were created in oil on canvas, with dimensions approximately corresponding to A3 size.

Subsequently, we enlisted two art-major students to replicate four paintings from each professional artist, resulting in a total of 20 forged pieces. The forgers were provided not only with the genuine paintings to be copied but also with detailed information regarding the paints used in the original works. This meticulous approach ensured that the forgeries were crafted with such precision that, to the untrained eye, distinguishing between the genuine paintings and the forgeries was nearly impossible at a cursory glance. Fig. 1 shows the authentic works and corresponding forgeries by artists a01 and a02.

To capture the entire artwork, we took high-resolution close-up photographs using a high-quality camera. We moved the camera at specific intervals in all directions, ensuring that the images slightly overlapped. During the post-processing stage, we cropped the images to ensure that they did not overlap. This resulted in approximately 78 images per artwork, each standardized to a size of 514x366 pixels and a resolution of 300dpi. With a total of 60 artworks, this process yielded over 3,800 images. Fig. 2 shows a set of 80 images capturing artist a02's artwork.

The final dataset consists of 50 paintings: 30 genuine works by professional artists and 20 forgeries by art students. For the purpose of training and testing the deep learning algorithms, we designated one genuine painting and its corresponding forgery from each artist as the test set, with the remaining works constituting the training set.

III. Proposed Network Structure

The architecture of the network we propose is divided into two main components: the feature extractor and the classifier. Fig. 3 shows a conceptual diagram of the proposed network, illustrating the integration of the feature extractor and classifier, and the unique utilization of dual image patch sets.

The feature extractor, as the name suggests, takes image patches as input and outputs a feature map. We utilized a pretrained ResNet module for this purpose. The rationale behind this choice stems from the limited size of our dataset. Despite having created the dataset ourselves, the small size posed a high risk of overfitting the data. By employing a ResNet module pretrained on the ImageNet dataset, we leveraged its ability to construct features from over a million images. This allowed for stable performance even with our relatively small dataset. To align with the pretrained ResNet-50 specifications, we randomly extracted patches of size 232x232 from our images.

It's worth noting that each image represents only a fraction of the entire artwork, capturing one of approximately 80 distinct regions. Extracting meaningful features from a single image patch to discern the characteristics of the entire artwork and determine the authenticity of the artist is a challenging task. This inherent difficulty underscores the need for our unique approach of using sets of image patches, which facilitates a more robust and comprehensive understanding of the artwork.

Unlike conventional methods that use a single image patch, our approach employs two sets of image patches. One set is extracted from images of the same artwork, while the other set is derived from images of a different artwork. If the size of the image patch set is given as 4, a total of 8 image patches are independently applied to the feature extractor. These multiple feature maps are then concatenated and passed to the classifier.

The classifier's size varies in proportion to the size of the image patch set. Its role is to classify whether the authors of the artworks from which the two image patch sets were extracted are the same or different. Consequently, our proposed network utilizes binary cross-entropy as the loss function.

IV. Experimental Results

We conducted artist-specific tests, including both authentic and forged pieces, with each artist contributing a total of 10 artworks. Two of these were allocated to the test set, and the remaining eight to the training set. We evaluated the accuracy of our proposed network by varying the size of the image patch sets, using sizes of 1, 2, and 4. The size of 1 served as a baseline, as it corresponds to the use of a single image patch.

Table I shows the experimental results. In the table, IPS-1, IPS-2, and IPS-4 refer to the network with image patch set sizes of 1, 2, and 4, respectively. The results reveal a clear trend of increasing accuracy with larger image set sizes across different datasets. For example, the accuracy for the a01 dataset improved from 92.01% with an image patch set size of 1 (IPS-1) to 99.88% with an image patch set size of 4 (IPS-4).

TABLE I: Experimental results of the proposed network

Artist Dataset	IPS-1	IPS-2	IPS-4
a01	92.01	98.37	**99.88**
a02	85.96	92.65	**96.13**
a03	79.89	84.57	**88.62**
a04	74.53	80.68	**86.72**
a05	72.60	72.84	**76.67**
Average	81.00	85.82	**89.60**

We can extrapolate from these results to predict the potential accuracy improvement with an image set size of 8. For a dataset like a01, where the accuracy is already high, the improvement may be marginal, possibly reaching close to 100%. In contrast, for a dataset like a05, where the accuracy is relatively lower, the improvement might be more substantial, potentially reaching around 80% or higher.

The use of a larger set of image patches proved to be more effective in representing the characteristics of the artwork, as evidenced by the consistent improvement in accuracy across different datasets. This supports our hypothesis that utilizing sets of image patches, rather than a single one, enhances the model's ability to discern the authenticity of the artist.

V. Conclusions

In this study, we introduced a novel deep learning approach to authenticate artworks by analyzing image patch sets. By employing a pretrained ResNet module as a feature extractor, we were able to leverage the knowledge gained from extensive pretraining on the ImageNet dataset, mitigating the risk of overfitting on our relatively small dataset.

Our unique approach of using sets of image patches, rather than single patches, proved to be effective. We demonstrated that increasing the size of the image patch set consistently improved the accuracy of the model across different artist datasets. The results, as summarized in the table, showed a substantial increase in accuracy from IPS-1 to IPS-4, with the most significant improvements observed in datasets with higher initial accuracy.

The creation of a custom dataset, consisting of both authentic and forged paintings, was an essential but challenging part of this research, enabling the intricate analysis required for the authentication task.

Future work may explore further optimizations, integration with other modalities, and application to a broader range of artistic styles and periods.

Acknowledgment

This research was supported by Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2020 (Project Name: Development of the system for digital data acquisition of modern and contemporary fine arts and supporting science-based art credibility analysis, Project Number: R2020060004)

REFERENCES

[1] F. Pozzi, J. Arslanoglu, and E. Nagy, "Alexander calder's half-circle, quarter-circle, and sphere (1932): a complex history of repainting unraveled," *Heritage Science*, vol. 8, no. 1, pp. 1–14, 2020.

[2] N. Van Noord, E. Hendriks, and E. Postma, "Toward discovery of the artist's style: Learning to recognize artists by their artworks," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 46–54, 2015.

[3] S. Hwang, H. Song, S.-W. Cho, C. E. Kim, C.-S. Kim, and K. Kim, "Optical measurements of paintings and the creation of an artwork database for authenticity," *Plos one*, vol. 12, no. 2, p. e0171354, 2017.

[4] A. Elgammal, Y. Kang, and M. Den Leeuw, "Picasso, matisse, or a fake? automated analysis of drawings at the stroke level for attribution and authentication," in *AAAI Conference on Artificial Intelligence*, 2018.

[5] A. Chen, R. Jesus, and M. Villarigues, "Using deep learning techniques for authentication of amadeo de souza cardoso paintings and drawings," in *EPIA Conference on Artificial Intelligence*. Springer, 2019, pp. 172–183.

[6] F. Ji, M. S. McMaster, S. Schwab, G. Singh, L. N. Smith, S. Adhikari, M. O'Dwyer, F. Sayed, A. Ingrisano, D. Yoder *et al.*, "Discerning the painter's hand: machine learning on surface topography," *Heritage Science*, vol. 9, no. 1, pp. 1–11, 2021.

[7] S. Kim, S. M. Park, S. Bak, G. H. Kim, C.-S. Kim, J. Jun, C. E. Kim, and K. Kim, "Investigation of craquelure patterns in oil paintings using precise 3d morphological analysis for art authentication," *Plos one*, vol. 17, no. 7, p. e0272078, 2022.