CLASSIFYING PAINTINGS OF AN ARTIST

Xiaoshan Liu, Danielle Williams

ABSTRACT

This project employs machine learning techniques to identify paintings by artists. We trained convolutional neural networks on a large dataset of paintings and compared the performance of two popular deep-learning architectures, ResNet and VGG, in terms of accuracy and training time. Our results showed that the model achieved high accuracy in identifying paintings by different artists, with both architectures performing similarly in accuracy but with differences in time. We also conducted experiments to investigate the effect of training dataset size on the model's accuracy and found that a smaller dataset could still achieve good accuracy as long as the artist's style is well-distinguished from others.

The ability to accurately identify the artist of a painting has several practical applications, such as historical research and education, and helps build online galleries. Our report provides a detailed description of the methods used, including data preprocessing, model architecture, and training and evaluation procedures. We also discuss the limitations and future directions of the project, including the potential biases in the dataset, the interpretability of the model, and the scalability of the approach to larger datasets or real-time applications.

1. INTRODUCTION

Throughout history, artists have employed various styles and techniques to create unique works of art that reflect their individuality and creativity. In the modern era, many artists have developed their own distinctive styles, allowing their works to be identified based on certain features. These unique styles have made these artists renowned as masters of their craft. In this project, we aimed to use machine learning techniques to identify paintings by several modern artists. Specifically, we trained and compared the performance of two popular deep learning architectures, ResNet and VGG, on a large dataset of paintings. Our goal was to predict the artist of a new painting based on its features, enabling us to automatically identify the work of a particular artist.

2. METHOD

Data Collection and Preprocessing. There are two datasets in our experiment. The first dataset from Kaggle[1] consists of 400 training images, 100 validation images, and 70 test images respectively for each of the 10 famous impressionist artists. The second dataset was also from Kaggle[6], we modified the original dataset and the new dataset had 50 training images and 20 validation images respectively for 40 artists from various time periods. Before training our model, we preprocessed the dataset by resizing all images to a fixed size of 256x256 pixels, and cropping the center part of the resized image to obtain a square of size 224x224 pixels.

Model Architecture. We applied pre-trained models directly. They were Resnet18, Resnet34, VGG16, VGG19.

Model Training and Evaluation. We trained the model using the SGD algorithm and the cross-entropy loss function. We trained the model for 10 epochs and monitored the validation accuracy to prevent overfitting. After training, we evaluated the performance of the model on the testing set using accuracy.

Experimental Setup. All experiments were conducted on Google Colab and using a standard GPU. We used several performance metrics, including accuracy and training time, to evaluate the effectiveness of the model and compare the performance of different CNN architectures.

3. EXPERIMENTS

We performed several experiments to evaluate the performance of different models trained on the same dataset. In our first comparison, we trained VGG16, VGG19, Resnet18, and Resnet34 and obtained accuracy values as shown in Table 1. All four models achieved high accuracy, with Resnet34 achieving the highest accuracy of 80.39%, while VGG16 had the lowest accuracy of 76.33%. Notably, Resnet models required significantly less training time than VGG models, with a more than 50% reduction in overall training time. This is consistent with the Resnet paper[2], as Resnet18 is a deeper network with fewer parameters per layer, thus yielding better performance and faster training time[2, 3, 4, 5].

	VGG16	VGG19	Resnet18	Resnet34
Accuracy On test data	76.33%	78.57%	79.97%	80.39%
Best Validation Accuracy	73.83%	76.56%	77.37%	77.68%
Training time	78m 57s	106m 1s	34m	38m
# of epochs	10	10	10	10

(Table 1)

In our second task, we investigated the number of images required to train a model effectively. We utilized a dataset consisting of 40 artists and trained a model on this smaller dataset. Our evaluation showed that approximately half of the artists achieved an accuracy of 60% or higher even with the reduced training dataset. This result suggests that the unique and prominent styles of these artists were accurately captured by the model despite the smaller training dataset. These findings offer insights into the role of data in training models and support the idea that some artists' styles are inherently distinct and can be learned effectively with fewer training images.

4. APPLICATIONS

Machine Learning has the potential to serve as a valuable reference tool for art historians and enthusiasts, helping to deepen our understanding of artistic styles. By analyzing misclassifications made by the model, we can gain insights into how artists learn and borrow from one another, and compare these findings with human interpretations of artistic influence.

Moreover, Machine Learning can be utilized to create online galleries, making it easier for collectors to discover and acquire works by their favorite artists. This technology can help to streamline the art acquisition process, providing collectors with a more efficient and personalized experience. These applications demonstrate the exciting possibilities that arise when Machine Learning and art intersect, and highlight the potential for this technology to enrich our understanding and appreciation of art.

5. LIMITATIONS

Here are several limitations to consider for this project.

Dataset bias. The accuracy of the model heavily relies on the quality and diversity of the training dataset. If the dataset is biased towards certain artists or styles, the model may not generalize well to new data. Moreover, if the dataset does not include enough examples of certain artists or styles, the model may not be able to accurately predict them.

Ambiguity in style. The style of an artwork is subjective and can be difficult to define objectively. This means that different experts may classify the same artwork differently, making it challenging to create a standardized dataset and train the model.

Lack of interpretability. Deep learning models can be difficult to interpret, making it challenging to understand why the model makes certain predictions. This can be problematic for applications that require a high degree of transparency and accountability.

Scalability. The training and evaluation of deep learning models can be computationally intensive and time-consuming, which can limit the scalability of the approach to larger datasets or real-time applications.

6. CONCLUSIONS

In this project, we used machine learning to identify paintings from several modern artists. We trained a model on a large dataset of paintings and used it to predict the artist of a new painting based on its features. Our results showed that the model achieved high accuracy in identifying paintings by different artists, demonstrating the potential of machine learning in the field of art analysis and creating online art galleries.

The success of our project highlights the importance of using advanced technologies to analyze and understand works of art. Our model can be used for a variety of applications, such as art authentication, historical research, and education. However, there are still limitations and challenges to overcome, such as the need for larger and more diverse datasets and the potential for biases in the training data.

Future research could focus on improving the model's performance, addressing the limitations and challenges, and expanding its applications. For example, it could be interesting to apply similar techniques to other types of artwork, such as sculpture or architecture, or to study the evolution of artistic styles over time. Overall, this project demonstrates the potential of machine learning in the field of art analysis and contributes to the ongoing efforts to bridge the gap between technology and the humanities.

7. REFERENCES

[1]A. Singh, “How to classify the paintings of an artist using Convolutional Neural Network,” Medium, 27-Oct-2020. [Online]. Available: https://abhisheksingh007226.medium.com/how-to-classify-the-paintings-of-an-artist-using-convolutional-neural-network-87e16c0b3ee0. [Accessed: 29-Apr-2023].

[2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.

[3] “[classic] deep residual learning for image recognition (paper explained),” YouTube, 14-Jul-2020. [Online]. Available:https://www.youtube.com/watch?v=GWt6Fu05voI. [Accessed: 29-Apr-2023].

[4] “[classic] deep residual learning for image recognition (paper explained),” YouTube, 14-Jul-2020. [Online]. Available:https://www.youtube.com/watch?v=GWt6Fu05voI. [Accessed: 29-Apr-2023].

[5] “Understanding VGG16: Concepts, architecture, and performance,” Datagen, 13-Apr-2023. [Online]. Available: https://datagen.tech/guides/computer-vision/vgg16/. [Accessed: 29-Apr-2023].

[6] Icaro, “Best artworks of All time,” Kaggle, 02-Mar-2019. [Online]. Available: https://www.kaggle.com/datasets/ikarus777/best-artworks-of-all-time. [Accessed: 29-Apr-2023].