ART RECOGNITION

AI Report by Art Recognition

AI authenticity evaluation of the artwork

«Samson and Delilah»

attributed to

Peter Paul Rubens

September 13, 2021

Digitally signed by Art Recognition AG
Date: 2021.09.26 22:05:25 +02'00'

Artwork Information

Attributed to	Peter Paul Rubens
Title	Samson and Delilah
Medium	Oil on Canvas
Image size	2263×2048 px
Training dataset	148 images

Summary

The Art Recognition AI System evaluates "Samson and Delilah" **NOT to be an original artwork** by Peter Paul Rubens with a probability of **91.78** % .

Method description

This report presents the results of an Artificial Intelligence analysis of the painting "Samson and Delilah" attributed to Flemish artist Peter Paul Rubens (1577–1640).
The purpose of this analysis is to determine wether the artist's main characteristics, learned by the AI algorithm from a dataset of original paintings by Peter Paul Rubens, match the characteristics identified on the painting in question. Below we describe the main steps of the process:

- The AI learns the artist's main features from a set of photographic reproductions of original paintings by Rubens. To improve the algorithm's discrimination capabilities, we also feed into the system a set of paintings which are not by Rubens.
- To increase precision, a series of detailed analyses are being carried out by splitting the training images into smaller patches and analysing each patch separately.
- Once the training has been completed, the AI compares the learned features with those identified on the painting "Samson and Delilah". Based on this comparison, the AI calculates a probability for the authenticity of the submitted artwork.

Keypoints of the analysis

1. Peter Paul Rubens is especially known for his biblical and mythological nudes. He combined the ideals of Italian renaissance and Flemish realism. Using many different colour tones, he arrived at realistic depitction of figures and nature. He mainly executed the preliminary work (sketches), whereas the actual realisation was then largely carried out by his workshop. His style remained more or less constant, with a bold and fast brushstroke, but in his last decade (1630-1640) his brushwork got looser and more tactile.
2. We trained the AI on a high quality image dataset containing original paintings fully created by Rubens (and not finalized by his workshop). The success of the training is confirmed by an accuracy on the validation set of 83.94%.
3. Following the successful training, the AI has assessed the authenticity of the image in question. The 91.78% negative class probability is a strong indicator of a mismatch between the brushstroke patterns learned from the training images, and those on the evaluated artwork.

   Finally, looking at the heat map, on the analysed painting the areas which most strongly support the algorithm's decision are Delilah's body as well as patches on the other figures (especially Samson's left arm and legs).

Detailed report

Data collection and preprocessing

We gathered a total of 148 original paintings by Peter Paul Rubens to train, optimize and validate our model. We have also fed into the system a set of contrast images which are not by Rubens, i.e., images of artworks by artists of similar styles and chronology.
Since the paintings exhibit high variability in terms of aspect-ratios, we employ a particular preprocessing strategy in order to capture both fine details and coarse structures. To capture finer details, we split the images into non-overlapping patches which are further cropped into squares and shuffled before being fed into the program. Patch sizes depend on the image quality and the size of the brushstrokes themselves. We also keep the full (undivided) images to capture the high-level structural elements.
The strategy described above allowed us to augment the original dataset, such that our final, balanced dataset contains a total of 2392 data stemming from original images and a similar number of data points in the contrast set. All images have been converted to the lossless Portable Network Graphics (PNG) format. The dataset generated as described above forms an excellent basis for training the AI.

Training the Algorithm

The deep convolutional neural network has firstly to learn the artist's discriminative features from the training data. The learning process is completely independent, in the sense that we do not introduce any hand-crafted features and thus do not not influence in any way the learning process. Aside from dividing the images into patches, we exploited several geometric transforms such as mirroring, rotation, and changes in scale and aspect-ratio of the input images. We also optimized the hyperparameters (learning rate, batch size, regularization parameters) to reach the most accurate model.
After completing the training, the robustness of the trained neural network has been tested on a validation set. Following a standard validation protocol, we split the images into training and testing set. This split was done randomly, preserving 80% of images for training, and 20% for testing. No painting in the training set appears in the testing set and vice-versa. The classification accuracy on the validation set is 83.94%.

Analysing the submitted image

Following the successful training, we passed the image of the painting "Samson and Delilah" through the AI, which compared the features learned from the training images with those on the image in question. Based on this comparison, the AI returned a class probability of 91.78% for a negative response ('not original').

Visualisation

In this section we present a heat map, which is meant to provide visual evidence on the algorithm's evaluation.

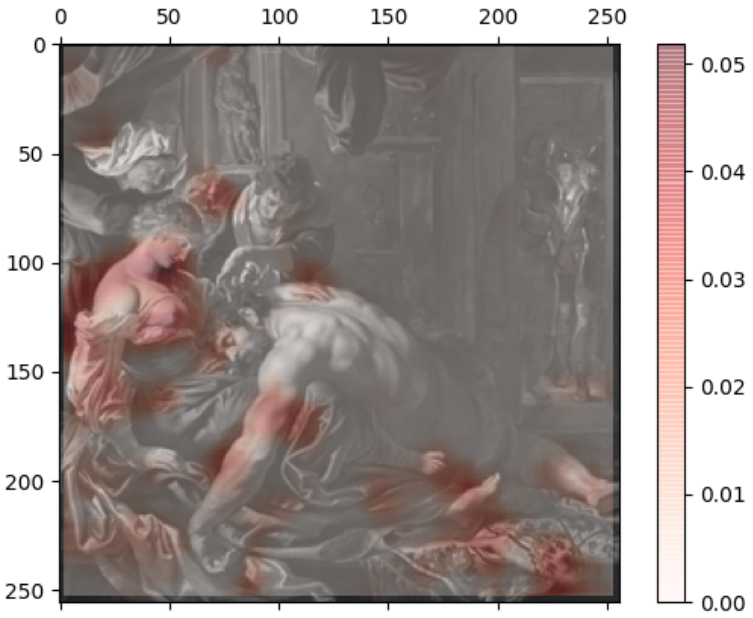

The interpretation is as follows: the most important areas for the algorithm's decision are those highlighted in red, and the importance gradually decreases proportional to the shade of red; at the other end of the spectrum, the regions which are not colored have the smallest influence on the decision. Therefore, the hotspot areas (colored most intensely in red) have the highest importance for the final decision.

The hotspots tend to appear in the regions which comprise more structure (corners, edges, shapes etc) or are important for the overall composition. But, they can also appear on flat surfaces where the brushstroke changes direction. An important observation is that these regions stem from the analysis of the readable brushstroke and other structural characteristics, and are not related to the artistic representation. Also note that the size and position of the regions marked in red are not directly linked to the class probability ('original'/'not original').

On the painting in question, the areas which most strongly support the algorithm's decision (not original) are Delilah's body as well as patches on the other figures (especially Samson's left arm and legs).

Short Glossary

ArtTech — an interdisciplinary domain combining art and technological innovation. As the art market continues to embrace new technologies, the aim is to break down barriers between art institutions, creative artists and technology startups, and offer new, digital services to art lovers and experts.

Unsupervised Deep Learning — a Machine Learning technique which uses Artificial Neural Networks to process large amounts of data. Deep Learning systems are able to learn features and patterns only based on data, and do not rely on rules programmed by a human. They are also able to make accurate decisions without the help of a human.

Deep Convolutional Neural Networks (DCNNs) — a class of Artificial Neural Networks consisting of many layers and nodes ("neurons") that process information hierarchically. The more layers and neurons a network has, the more complex it becomes (hence the attribute "deep"). Each layer can be understood as a collection of filters which extract certain feature from the input data. For example, in the first layer the DCNN can recognize simple structures such as lines or edges. The following layer can recognize combinations of those structures such as simple shapes or curves. With each layer the identified structures become increasingly complicated. In the last step (classification layer), the final results are assigned to a certain class.

DCNNs and Image Recognition. DCNNs are most powerful in image processing, and already in the early 2000 have been applied to the detection of objects and regions in images. Meanwhile DCNNs are the most dominant approach for almost all recognition tasks, thanks to a new regularization technique called dropout and techniques to generate more training examples by deforming the existing ones. The performance of CNNs has determined major technology companies (such as Google, Facebook, Microsoft) as well as a growing number of start-ups to initiate research projects and deploy DCNN-based products and services.

Training Data — a set of pre-labeled data (for instance, images labeled as 'original' or 'fake') which are fed into a Neural Network to learn (and subsequently recognize) particular patterns and features. In order to have a network performing very good on unlabeled data after completing the training, it is important to ensure that the input data is of excellent quality from trustworthy sources.

Validation Protocol — a set of established standards to assess the performance of a learned AI model. These include using independent training and testing sets, and making sure that feature selection is performed using the training set only. In the case of small datasets, the model perfomance can also be evaluated by using cross-validation techniques. A standard performance measure is the validation accuracy.

Model Interpretability — refers to how easy it is for humans to understand the neural networks' decision process and outcome. Until recently, AI algorithms have been notoriously "black boxes," providing no way to understand their inner processes and making it difficult to explain the insights even to experts. Meanwhile this lack of transparency can be overcome with the help of heat maps which expose which features are most impactful to the model and how changes in the values of each feature affect the overall outcomes.

Heat Map — is an image which visualizes the importance of each pixel for the prediction. This type of images are very useful to understand why the AI System has arrived at this decision, and check wether they match human intuition.

Neural Style Transfer — an AI technique that blends the "content" (objects and their arrangement in the image) with the "style" (texture of an image). This is particularly useful for identifying the brushstroke of an artist.