

Discovering ghost paintings using AI and neural networks and its importance in the art world

Ani A. Parnagian – Fall 2021

IPHS 200 – Programming Humanity, Professors Chun and Elkins

Background

For most of art restoration history, x-radiography and infrared reflectography imaging was used to look underneath the paint: to understand damage, but also to see if there were underdrawings (preparatory sketches) or ghost paintings (full paintings that were covered over). Even with this technology, there were big problems. The visible painting and the ghost paintings lines would be combined when it was x-rayed, and it means that a human would have to do the painstaking work of separating the two layers. Additionally, while x-rays could tell what the material was (different materials expel different x-rays), they could not tell what color was used by the artist on the ghost painting.

The author of these papers use deep neural networks and utilize many images from the artists other works to match the “style” of the piece underneath. They use the deep neural networks to generate full color realistic images that consider many factors of the artist’s work to produce a work that would otherwise have been completely lost. What is so special about this technology is that it “combines the power of deep learning with properties such as special invariance which is inherent to visual patterns.” In other words, it utilizes a very human way of looking at negative space, and it combines it with the mathematical precision and numerous style images that only a computer can process. This technology is new and these are some of the earliest examples of using it for art restoration and curation.

Technology Background

This project utilizes a couple different types of recently developed technology:

Convolutional neural networks: these neural networks are most often applied to analyze visual imagery, and for this project are particularly good at image segmentation, needed to divide the top image from the bottom image. The input image is then to a filter that uses weights that measure if the pixels taken from the original are close to what we are trying to find in an image (i.e.. the woman underneath Picasso's *Old Guitarist*) when we are using it for identification and identifies the smaller sections of the figure as a whole.

Deep neural networks: a deep neural network is used to describe a neural network with more than one layer, allowing for more complex ways to process data and utilizing hidden networks, which store information and to most of the work using weights. Deep neural networks allow the machine to generalize on their own and make effective predictions. The deep neural networks in this project do this to create the best predictions possible of the paintings hidden underneath.

SSSSL: semi-semantic, semi-supervised labelling is a way to label a dataset using the small amount of labeled data to understand the rest of the unlabeled data. This technology uses human labeling of regions of an image (for example, in a portrait, labeling a hand or a foot) and once the machine has that information on some of the sections, it can generalize it to the rest of the image.

Methodology

The first step in recovering a lost image is image separation of the primary visible image from the ghost-painting underneath. This is done by weighting the images and splitting the value of each individual pixel between the two images, so that it matches it creates two clearly separate images but shows cohesive lines. This is particularly difficult with painting because the images are not highly linear, and the rules of painting are much more varied than the much more ordered art of photography, and comparable images are much harder to find for older paintings, but relatively easy for photographs.

Then, neural style transfer is used. In some cases, such as with the painting of Pablo Picasso’s *Old Guitarist*, only on image is used to add color, and contour. And this blends them together, using the linework of the “content image” (the outline of the ghost painting) and combines it with the style image: *La Vie*, another piece by Picasso during his blue period, and creates the final-colored image of the ghost painting. Historical information, like when this piece was painted in Picasso’s career, is used to help find the most probable color palate for the ghost painting.

This is done using the VGG-Network, a neural network. The first layer of the network is basic information that is fed in, like the style image and the content image. As the network continues, they look for more specific features, and take information from both the style and content image, breaking then down into individual parts to understand each color and each line, and can understand how these pixels create the images they were given.

This process can also be done with an ensemble of style images, because usually one style image does not represent the style of an artist or the period. What is most important in this type of transfer is transfer of form a color, and this is completed using bounding contours to complete the lines. In the case of Leonardo Da Vinci’s *Virgin of the Rocks*, it uses edge maps of Leonardo’s other paintings to fill in lines that the x-ray does not pick up on while preserving Leonardo’s style.

To enhance resolution of these newly discovered paintings, there is even more style transfer used by utilizing generative adversarial deep neural networks, or GANs. This type of neural network uses two sub-networks: the generator network that generates new images, and the discriminator network that classifies by style and makes decisions, fine tuning the resolution of an image pixel by pixel. The GAN learned how to take low resolution images to high resolution images by comparing low resolution and high-resolution image of Leonardo’s work, and the sub-networks described above created a map of the difference between the low and high quality, so it is able to do so successfully when a new piece is given to the system.

Even after it goes through the system, there is fine tuning that needs to be done, and colors seem to bleed into each other, and textures are not sharp. While many of these problems could be resolved with more style images, there are often not many, especially with older pieces, and because more images were not available, it had to be done partly by hand, using semi-semantic, semi-supervised labelling (SSSSL) to mark specific areas on a body or a landscape, making the final image much better.

While the AI renderings do not look perfect yet, this is a big step in digital art historical analysis and will have a big effect on the way art historians, conservators and curators work.




Figure 1. Works by Picasso used in the computational estimation of the ghost-paintings in *Old guitarist*. a) *Old guitarist* (122.9 × 82.6 cm), oil on canvas (1903-04), Art Institute of Chicago, b) x-ray of *Old guitarist*, Art Institute of Chicago, c) the “content image,” i.e., the hand-edited underdrawing of *Old guitarist* in grayscale, d) the “style image,” *La vie* (196.5 × 129.2 cm), oil on canvas (1903), Cleveland Museum of Art, and e) the computed portrait, where the color style from d) has been transferred to the design from c) by means of deep neural networks.



Figure 4. An overview of our method for extraction of the ghost-painting in Leonardo’s *Virgin of the rocks* and style transfer based on an ensemble of Leonardo’s works. At the left is an x-ray of *Virgin of the rocks*, and next is a color-inverted version of the x-ray, which facilitates incorporation of Leonardo’s traditional drawings and artworks (dark lines on white support). Next is a set of representative drawings by the artist, each of which leads, through style transfer, to a corresponding contour. These are then integrated to yield the final, contour-completed model of the underdrawing, at the right.




Figure 5. a) Leonardo’s *Virgin of the rocks* (189.5 × 120 cm), oil on panel (1495-1508), National Gallery London. b) The grayscale underdrawing of a Madonna and child composition, revealed by x-radiography. c) Human-generated coarse segmentation of the underdrawing. d) A detail showing just the Virgin’s head from panel c). e) The computed full-color ghost-painting of the passage in d).




Figure 7. a) The full inferred ghost-painting behind Leonardo’s *Virgin of the rocks*, computed by our convolutional generative adversarial networks. Clearly the style does not fully represent that of Leonardo: the contours are rather sinuous and the skin modeling is uneven and mottled, unlike the sfumato and chiaroscuro that characterizes his portraits. b) A set of 16 inferred works used to derive a).

Uses and Ethics

There is not much to consider in the negative aspects of this technology, but it is important to consider the ethics of uncovering artwork that the artists did not mean to be seen or explored as part of their work. Is it ethical to uncover this work, and if this technology is ever applied to artists who are still alive, would this change the understanding of ethics of hidden works? Additionally, the company that the authors of these papers work on sell their AI creations on their website, ranging as of now from \$5,000 to \$20,000. Is it ethical for them to sell these works, and would it be considered their creation because they used the technology to uncover it? This however could have a positive effect as it allows artwork by famous artists to be accessible to many more people.

In terms of positive effects and how this will expand the art world, there are numerous ways that this technology can be utilized. I believe that this technology will allow for job growth, as museums and galleries realize the artwork that is underneath existing paintings and how it can make the piece much more valuable and a greater investment, therefore expanding the art market. This will mean that more computer scientists will be involved in the art world, which will likely foster more involvement of tech in the art world, exposing more ghost paintings and getting a greater understanding of what the process was of these famous painters. It additionally allows for newer styles of artwork, like NFTs, to become a greater part of the art world with older paintings, not just with new pieces.

Additionally, this will allow art historical academic studies to grow, allowing both professors and students to better understand the oeuvre of the artists, especially because pieces that have underdrawings are often from the beginning of an artists career, where historians often have less information – in other words, it will help researchers better understand what they are studying and how an artist or a style evolved over time, giving humans a better understanding of our own artistic history.

As for what the developers of this technology hope, they say that their ultimate goal is “to someday recover hidden works with such fidelity to the original artistic intention that the computed works are worthy of the careful connoisseurial analysis enjoyed by our ‘visible’ cultural patrimony.”

Citations

Bourached, Anthony & Cann, George. (2019). Raiders of the Lost Art.

Bourached, Anthony & Cann, George, & Griffiths, Ryan-Rhys & G. Stork, David. (2021). Recovery of underdrawings and ghost paintings via style transfer by deep convolutional neural networks: A digital tool for art scholars.

Bourached, Anthony & Cann, George, & Griffiths, Ryan-Rhys & G. Stork, David. (2021). Resolution enhancement in the recovery of underdrawings via style transfer by generative adversarial deep neural networks.

www.oxia-palus.com

Acknowledgements

I would like to thank Professor Chun and Professor Elkins for their guidance through this course and this project in understanding AI technology and deep neural networks, as well as guiding me through the ethical dilemmas surrounding this new technology.