# Artistic Influence GAN

**Eric Chu**
MIT Media Lab
echu@mit.edu

## Abstract

What if Banksy had met Jackson Pollock during his formative years, or if David Hockney had missed out on the Tate Gallery's famous 1960 Picasso exhibition? How would their subsequent art differ? Inspired by these "what if" questions around artistic influence, we collected a dataset of paintings and the directed links of influence between artists. We then introduced an Artistic Influence generative adversarial network (GAN), in which the generator takes as input not only the noise vector $\mathbf{z}$, but also an additional embedding $\mathbf{v}$ representing the influencers. Thus, at inference time, we can imagine a new artist $A$ by specifying which artists influenced that artist $A$, and use the generator to produce paintings.

## 1 Introduction

Every piece of art is contextualized by the artist's background and the time period in which they lived, and part of this contextualization includes understanding which other artists influenced that artist.

Previous research has attempted to automatically discover these links of artistic influence by detecting common thematic and stylistic traces through machine learning [1]. There has also been many recent developments that use machine learning to generate art, ranging from transferring the style of a painting to a natural image [2] to generating paintings using generative adversarial networks [3].

We attempted to combine these two threads of research to further examine the question of how influence affects the work produced, by (a) collecting a dataset of images of art, with information about the influence links between artists, and (b) training a network that generates paintings conditioned on an "influence" embedding. In our final system, the user can "create" a new artist $A$ by specifying which artists influenced $A$. Images would then be generated as if they were painted by that new artist.

The dataset was collected from Wikiart[1], with a total of over 121,405 paintings from 2,539 artists. The website also provides metadata about each artist, including which artists influenced whom (e.g. Karl Otto Gotz was influenced by Max Ernst, Juan Gris, Wassily Kandinsky, and Paul Klee). After filtering the dataset to only those artists with known influence links, the final dataset is comprised of 27,138 paintings from 202 artists.

## 2 Model

### 2.1 Painting Embeddings

We represent a painting by extracting content and style using a pre-trained VGG object recognition network [4]. The content vector is the penultimate layer before classification, while the style vector is represented using a Gram matrix of the separate convolutional filters within a layer. This representation has been useful for both style transfer of images [2] and style-based image retrieval [5]. The style vector is created using the 4th, 5th, and 6th layers of the network, and we also use PCA for dimensionality reduction following [5].

[1] `https://www.wikiart.org/`

Submitted to 32nd Conference on Neural Information Processing Systems (NIPS 2018). Do not distribute.

### 2.2 Artist Embeddings

An artist could be naively represented as the average of all his or her painting embeddings. However, artists often have periods of varying styles over the course of their life. To capture this multi-modal distribution, we represent each artist as the means of a Gaussian mixture model (GMM). To select an appropriate number of components, we perform hyper-parameter search over the number of components and the covariance structure for each artist. Across all artists, the median number of components was 2, and the average number of components was 3.18.

### 2.3 Artistic Influence GAN

We use a deep convolutional generative adversarial network. In addition to the noise vector $\boldsymbol{z}$, the generator also takes as input an "influence" embedding $\boldsymbol{v}$. For a real image $x$ produced by artist $a^{(i)}$, we calculate $\boldsymbol{v}$ using the artist embeddings for all the artists that influenced $a^{(i)}$. We tested two variants for computing $\boldsymbol{v}$. The first simply averages the GMM components across influencers. The second passes all the GMM components into a LSTM (sorting the influencers by birth year in order to potentially account for the temporal effects), and outputs the final hidden state as $\boldsymbol{z}$. We found the LSTM model to sometimes produce qualitatively more interesting results, but also be more difficult to optimize.

The generator has six transposed convolution layers with batch normalization [7] and leaky ReLUs [8] and produces a final image of size 128. The discriminator has five convolution plus batch normalization layers. We use two losses to train our model:

1. Adversarial loss as described in the original GAN work [9], which drives the generator to produces images that look like paintings in the dataset.
2. Auxiliary classification loss as described in [10] to predict the artist of the painting.

## 3 Results and Discussion

Examples are shown in Figure 1. We find that when only one influencer is specified, the model produces paintings that look similar in style to paintings by that artist. For example, we find the results to resemble the broad vertical strips of color by Morris Louis, the landscape paintings by William Turner, and the kaleidoscope arrays by Paul Klee. When we specify the influencers as a landscape painter such as Turner or Thomas Cole plus an Expressionist like Klee, we find the muted, brown tones typically found in landscape paintings infused with a splash of color. However, the model struggles when more influencers are specified, with the images typically being some form of gray-green noise.

We believe that the quality of the generated images is greatly limited by the relatively standard GAN architecture we used. The model could benefit greatly by incorporating recent advances in GAN architecture and losses in order to produce larger and more "accurate" images. The model could also be tweaked to incorporate *influencees*. This would allow users to specify two artists (with distinct styles) and allow the model to generate paintings that would bridge the two. Ultimately, we hope aspiring artists or art students could use a tool built on top of this model to examine pockets of art history and find inspiration for their own art.

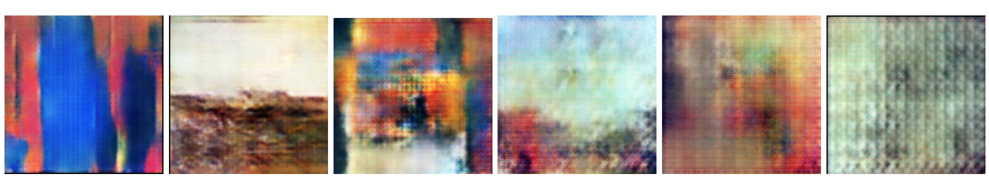

(a) Morris Louis (b) William Turner (c) Paul Klee (d) Paul Klee and William Turner (e) Judy Chicago and Thomas Cole (f) Caravaggio, de Kooning, Raphael, and Picasso

Figure 1: Examples from Artistic Influence GAN, with the user-specified influencers used to generate the painting noted below each image.

## References

[1] Saleh, Babak, et al. "Toward automated discovery of artistic influence." Multimedia Tools and Applications 75.7 (2016): 3565-3591.

[2] Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "A neural algorithm of artistic style." arXiv preprint arXiv:1508.06576 (2015).

[3] Tan, Wei Ren, et al. "ArtGAN: Artwork synthesis with conditional categorical GANs." Image Processing (ICIP), 2017 IEEE International Conference on. IEEE, 2017.

[4] Simonyan, Karen, and Andrew Zisserman. "Very deep convolutional networks for large-scale image recognition." arXiv preprint arXiv:1409.1556 (2014).

[5] Matsuo, Shin, and Keiji Yanai. "Style Image Retrieval Using CNN-based Style Vector."

[6] Radford, Alec, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks." arXiv preprint arXiv:1511.06434 (2015).

[7] Ioffe, Sergey, and Christian Szegedy. "Batch normalization: Accelerating deep network training by reducing internal covariate shift." arXiv preprint arXiv:1502.03167 (2015).

[8] Maas, Andrew L., Awni Y. Hannun, and Andrew Y. Ng. "Rectifier nonlinearities improve neural network acoustic models." Proc. icml. Vol. 30. No. 1. 2013.

[9] Goodfellow, Ian, et al. "Generative adversarial nets." Advances in neural information processing systems. 2014.

[10] Odena, Augustus, Christopher Olah, and Jonathon Shlens. "Conditional image synthesis with auxiliary classifier gans." arXiv preprint arXiv:1610.09585 (2016).