

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 with metadata on 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. At inference time, we can then imagine a new artist \mathcal{A} by specifying which artists influenced that artist \mathcal{A} , and use the generator to produce paintings.

PREPROCESSING

Data

- Wikiart: 121,405 paintings from 2,539 artists.
- Metadata includes which artist influenced whom (e.g. Morris Louis was influenced by Helen Frankenthaler, Joan Miro, and Jackson Pollock).
- After filtering the dataset to artists with known influence links: 27,138 paintings from 202 artists.

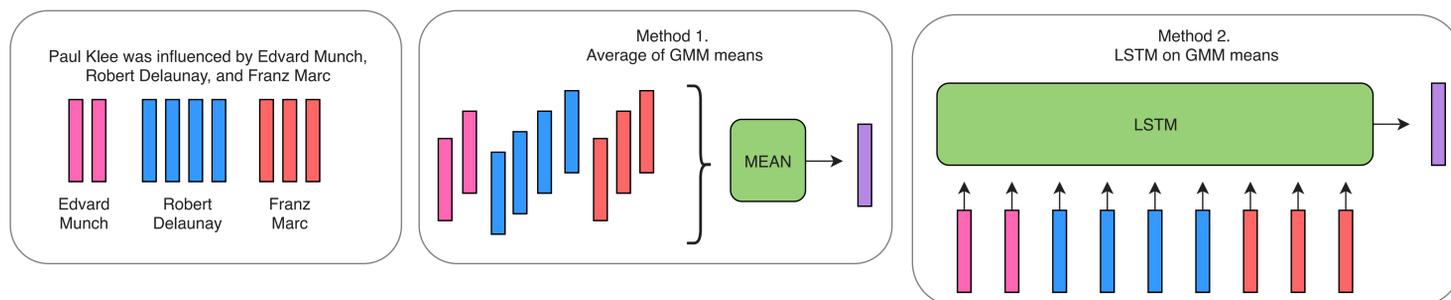
Painting Embedding

- Extract content and style vectors from pretrained VGG object recognition network
- Content vector: activation of penultimate layer.
- Style vector: Gram matrix of the separate convolutional filters within a layer.
- We also use PCA for dimensionality reduction.

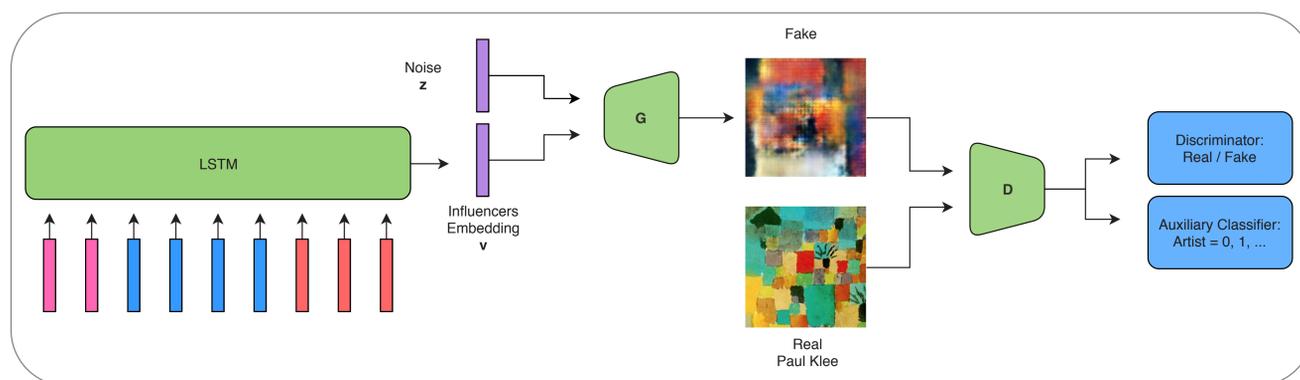
Artist Embedding

- An artist could be naively represented as the average of their 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) over their painting embeddings.**
- Across all artists, the median number of components was 2, and the mean number of components was 3.18.

INFLUENCERS EMBEDDING

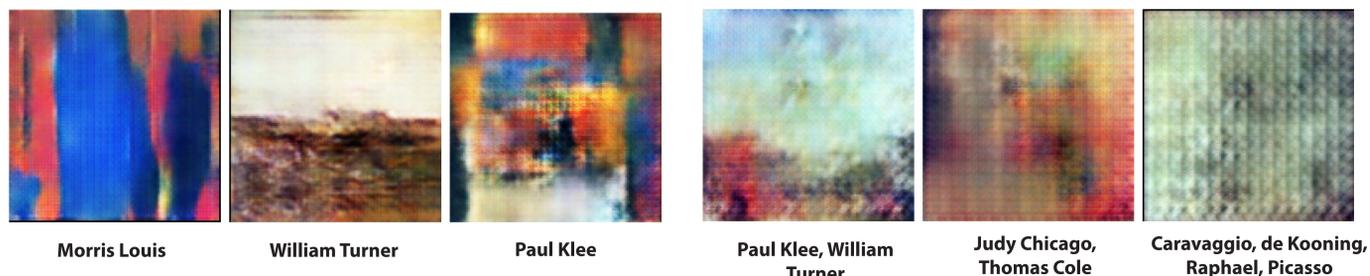


ARTISTIC INFLUENCE GAN



RESULTS

Influencers specified at inference time are listed below each generated image



Discussion

- 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.

Limitations and Future Possibilities

- Poor results when 3 or more influencers are specified, with the images typically being some form of gray-green noise.
- We believe the quality is limited by the relatively standard GAN architecture we used (DCGAN from 2015). The model could benefit greatly from recent advances in GAN architecture and losses.
- 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