10-615 Art and Machine Learning

Art for All:

Project One

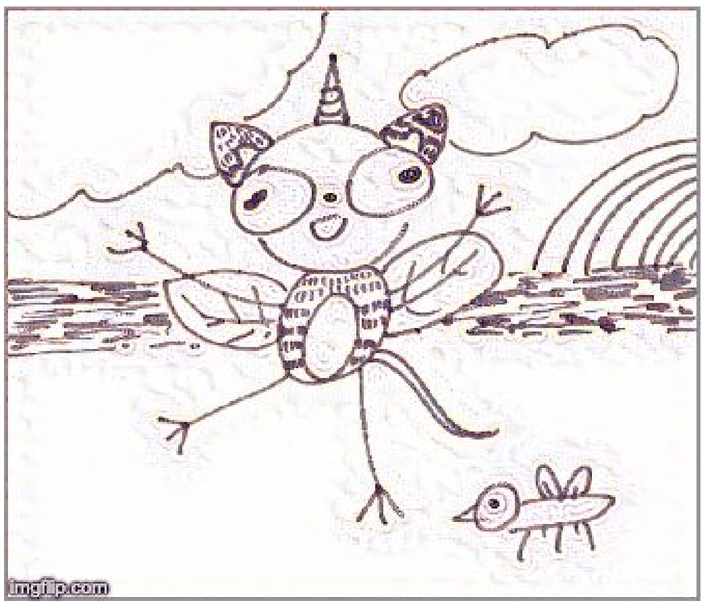

Kevin Song

Stuart McGibbon

Stella Kim

Tejas J. Prakash

Angel C. Hernandez

Concept

The concept at the center of our project was "the democratization of art." In short, we wanted to explore the possibility of whether the style transfer method could be used to turn laypeople's hand sketches into pieces of art resembling the works of their favorite visual artists. It was our hope that our project might empower individuals with little confidence in their creative abilities, by demonstrating the promise of human-machine co-creation. To collect the material for our project, we distributed a survey to students around CMU asking them to provide a sketch and a list of their favorite artists and styles. These responses provided us with the content and style images that would be the basis of our experimentation.

Technique

At a high-level, style transfer takes a style image and a content image and tries to map an optimal amount of style from the style image to the content image. Optimal is defined by two functions: the lstyle and lcontent which describe the numerical difference between style and content of two images, respectively. Given a third image, the output, we use backpropagation to broadcast the style and content onto the output image while minimizing the total of the lstyle and lcontent functions. The features used in the algorithm are generated from the hidden layers of pre-trained convolutional neural network, VGG19, which was trained to classify images. We generate features from this network because each layer of the network is designed to capture some nuance in a given image. The source notebook we used in when conducting this project already had loss functions, backpropagation and the VGG19 network loaded. As result, our team was concerned with editing the hyperparameters to produce aesthetic images over our hand-drawn dataset. The hyper parameters we tuned are as followed:

- **Iterations:** the number of epochs used to generate the output image
- **Style weight:** the amount of weight allocated to style in the loss functions
- **Content weight:** the amount of weight allocated to content in the loss functions
- **Content features:** the hidden layer(s) used in VGG19 to generate the content features
- **Style features:** the hidden layers(s) used in VGG19 to generate the content features

The methods used to tune these parameters were explicit enumeration[1] and randomly selecting the hidden layers within VGG19 used for generating the features, over a uniform distribution. Furthermore, things like recursively transferring style to a content image and adding random noise to the image were tested, as well.

Process and Result

The following section summarizes the success and fail rates of different methodologies used to assign values to the hyper parameters.

Failures Using Explicit Enumeration

The following images were generated using the recommended settings configuration of hyperparameters with some slight variations. At this point, the team was still in an exploratory
phase and felt it was best to begin in a conservative environment. Unfortunately, we quickly released the source of our content images required more hyperparameter tuning.

[1] Explicit enumeration: manually assigning different values to parameters and using results to drive decision making in future trials

Figure 1

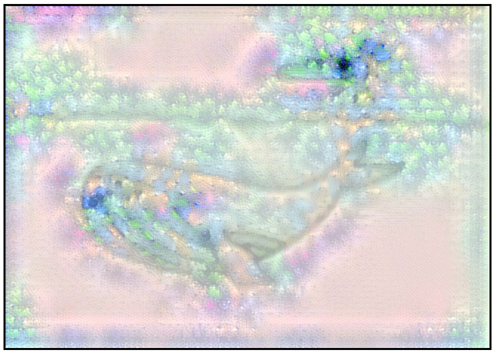

Figure 2

Here we can clearly see the style was not cohesive with the content image. Style and content images found in Appendix

Ultimately, we concluded that we had too much white space in our drawings since they were simple doodles with sharpie on a white paper, and there was not enough 'style' that the AI could catch and transfer. As result, we continued to explore different hyperparameter settings that deviated from the recommended environment with hope of the model picking up on different nuances in the images.

Success Using Recursion and Random Parameters

Figure 3 was generated by feeding the output as a new input and running the style transfer recursively and we notice the transferred style more evidently. Furthermore, **Figure 4** was generated by randoming selecting the hidden layers that were used to generate content and style features and we notice a clean abstract image of chair.

Figure 3

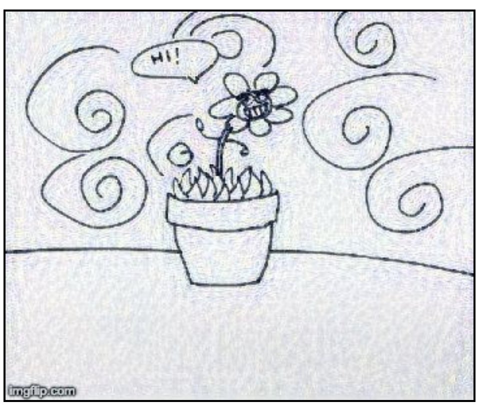

Figure 4

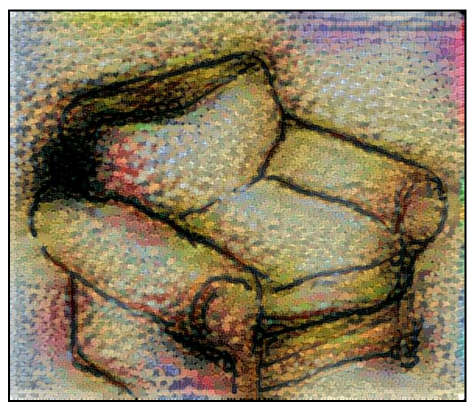

Content and style images can be found in Appendix.

During these trial runs we learned the algorithm could generate aesthetic images with the proper settings and little ingenuity.

Success Injecting Random Noise

In these trials we used the a [web app] to inject random noise to our images. We felt this would give the algorithm more nuances to catch onto and felt this was a good replacement of the solid white background across all our hand drawn images.

Figure 5

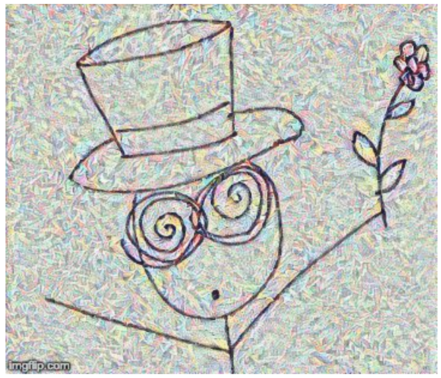

Style and content can be found in Appendix

Here we notice a nice collage of colors which was made possible by injecting random noise. You can find many of our other images and trials at our Git repo!

Code: https://github.com/ahernandez105/artml_project1

Reflection

At the outset of the project, our team set the rather ambitious goal of "democratizing art," by way of turning laypeople's sketches into pieces of art resembling the works of their favorite visual artists. Along the way, we learned a great deal about our chosen method (style transfer), including its strengths and limitations. For example, the process of transferring a chosen artist's style to the relatively minimal sketches (originally done in sharpie on paper) revealed certain limitations about the style transfer method, and fundamental challenges of the "democratization" concept, as we interpreted it. In particular, because we had the opportunity to apply so many artists' styles to the sketches, we came to understand that the style transfer method works best with content by artists who have a bold"visual style" -- i.e. highly textured, visually salient, often abstract -- rather than those that tend to lean "realist." Moreover, certain sketches seemed to lend themselves to stronger outcomes than others. Though some of our participants opted to draw full scenes, landscapes, or human portraits, experimentation revealed that the "random" nature of the style transfer method generated results that were less aesthetically appealing and more difficult to interpret when the content image was a complex scene. Rather, simpler representations such as the still life bottle, chair, and flower vase yielded surprisingly beautiful outcomes.

In summary, it was perhaps a bit ambitious to think we could turn anyone's drawing into a work resembling their favorite artists, particularly given the constraints on time and resources. However, our project produced some undeniably pleasing results and raised some interesting questions in the process. In the end, we acknowledge that "the democratization of art" is still a long way off, but hope that our project has stoked interest in the possibility of co-creation between humans and machines, and we hope that others will continue to explore this idea.

Each Member's Contribution

Stuart -

Stella - generating images, creating gifs, report and presentation slides

Angel - code edits, generating images, git repo and report

Kevin - generating images, report, and presentation slides

Tejas - presentation, paper

Appendix

Figure 1 Content and Style Images

Jing.png

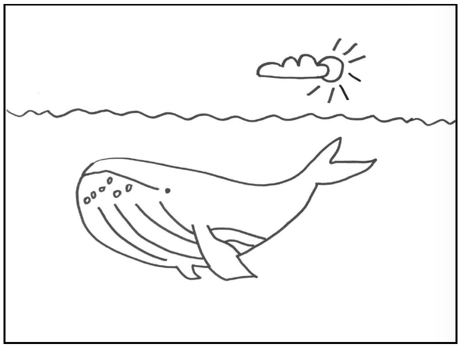

Beach in Pourville, Claude Monet

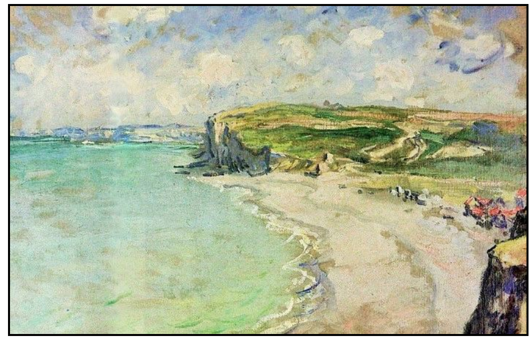

Figure 2 Content and Style Images

Gaby.png

Tim Burton

Figure 3 Content and Style Images

Jihoon.png

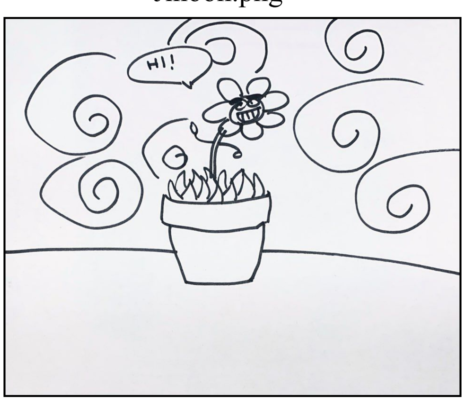

Van Gogh

Figure 4 Content and Style Images

young.png

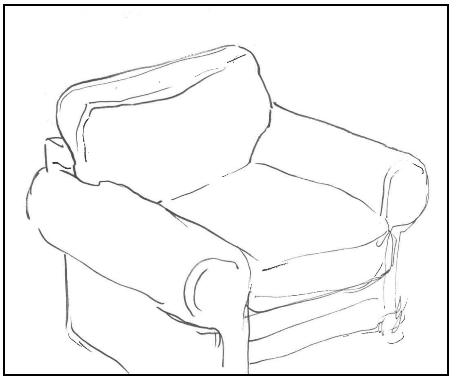

Garden at Vaucresson
Édouard Vuillard

Figure 5 Content and Style Images

kevin.png

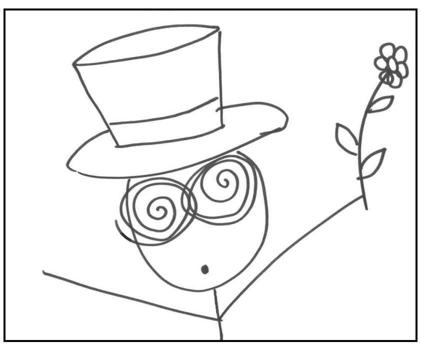

kevin.png with noise

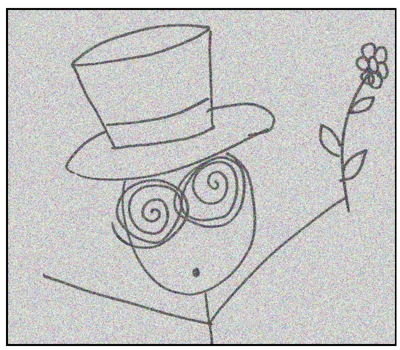

Salvador Dali

Figure 6 Content , Style, Output Images (Iteration=20, style=10)

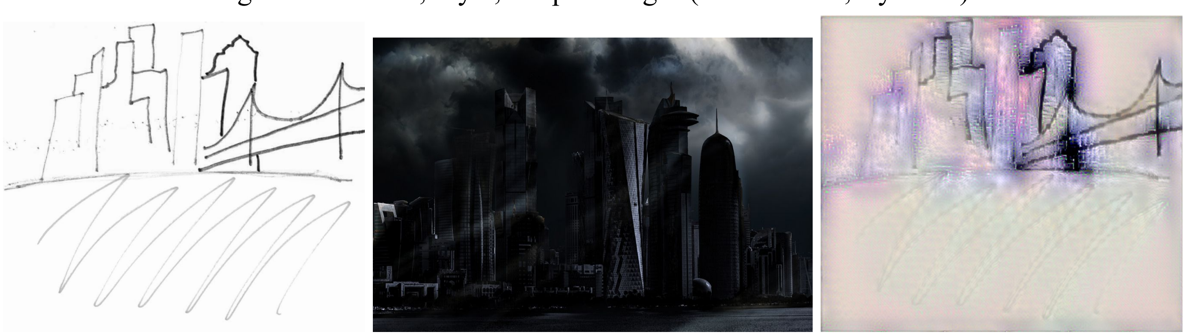

Figure 7: Selected Results

7a: Sketch by Stella Kim, Style Image by Stanley Donwood

7b: Sketch by Gillis Bernard, Style Image by Edgar Degas

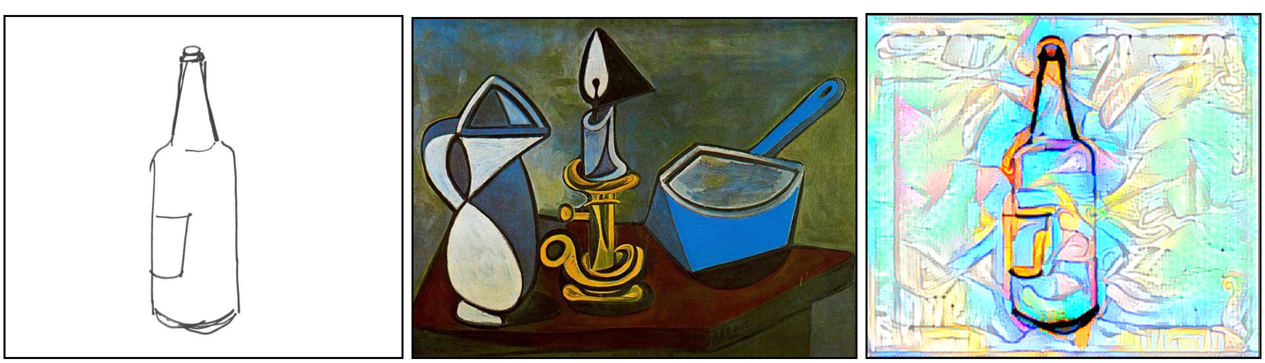

7c: Sketch by Elva Fu, Style Image by Pablo Picasso

7d: Sketch by Advita, Style Image by Gustav Klimt