

CS 523 S2017
Multimedia Systems
Jan 24, 2017

Project 1 – Implement a Generative ML Algorithm for Processing Image Data

This assignment may be completed individually or in teams of two or three.

Part 1 – due 1/31

Part 2 – due 2/7

Choose one of the two following projects to implement using the TensorFlow library:

A.) Image Classification + DeepDream

- Create a simple image classifier for detecting if a feature is present in an image or not. You can use the Inception v3 network (provided by Google), retraining it to recognize a new class of images, or you can build your own (maybe based off one of the TensorFlow MNIST tutorials, but with a custom dataset to classify.)

- Implement the Deep Dream algorithm for “hallucinating” features. Experiment with different parameters to explore different kinds of effects. Create an interactive interface that lets you input new photos and tweak their output.

<https://codelabs.developers.google.com/codelabs/tensorflow-for-poets/>

https://www.tensorflow.org/tutorials/image_recognition/

<https://www.youtube.com/watch?v=MrBzgvUNr4w>

<https://github.com/tensorflow/tensorflow/blob/master/tensorflow/examples/tutorials/deepdream/deepdream.ipynb>

B.) Implement NeuralStyle

- Implement a version of the original NeuralStyle algorithm by Gatys et al., or one of the newer algorithms inspired by it. Train the algorithm to generate at least one new style. Create an interactive interface that lets you input new photos and stylize them, and to explore different stylization parameters.

<https://github.com/lengstrom/fast-style-transfer>

<https://github.com/anishathalye/neural-style>

<https://arxiv.org/pdf/1508.06576v2.pdf>

(At least three groups must focus on A, and three on B.)

For a grade of C:

You will successfully implement example code provided in an online tutorial or taken from a code repository, etc. You will clearly document the code and explain what each line of code is doing (especially the code related to training or sampling the neural network). You will put your project on a GitHub repo with clear instructions for how to install and run and a clear explanation for what it is doing.

For a grade of B:

Additionally, you will create an interactive interface makes it easy for a user to upload a photo and to tweak parameters to generate new output.

For a grade of A:

Additionally, you will generate novel content based on creating a new style (for the style transfer) or new image classifications (for Deep Dream), or in some other way demonstrate understanding of the code by “making it your own.”

Part 1, due 1/31 in class

- Choose a paper, tutorial, or existing repo from which to base your implementation. (That is, get some code to run.)**
- Explain some technical details of your implementation: What type of neural network does it use? How many hidden layers does it have? How does it learn or make use of image features? (That is, explain the code.)**

Part 2, due 2/7 in class

- Create an interactive interface (in a Jupyter notebook, using a HTML/Javascript library, or via a Python UI) that lets a user upload a photo and tweak parameters to create a novel output.**
- Present your work in class, and explain how you implemented and extended previous work.**
- Describe what you have learned about neural networks.**