CSCI 1051 Project

Submission Instructions

Please upload your proposal to Canvas by **5pm Monday January 23, 2023**.

Project Proposal (6 points)

Your course project (worth 30 points total) is an opportunity for you to explore a deep learning topic that is interesting to you. You can work on the project by yourself or with a partner. The purpose of the proposal is to get you thinking about your project.

Unfortunately, we will not be able to cover every topic in the class before you'll have to choose what you want to do. So I have made a list of some potential projects that I will explain them to you in class. If you have questions, please ask via Canvas discussion or in office hours.

Project Ideas

The first step is to figure out the topic you're interested in. I have curated the following "wish list" of topics I'd like to learn more about:

- *Neural style transfer*: given a 'content' image and 'style' image, make a version of the 'content' image in the style of the 'style' image
- Diffusion from scratch on a small data set: generate meaningful images from noise (this is one of the key techniques behind stable diffusion)
- *Variational autoencoder*: train a meaningful embedding spaces for images and generate *new* images with random vectors in the embedding space
- *Adversarial input to image classification*: given an image classification model, find examples of images that look like one object to humans but appear as a completely different object to the classification image
- *Class activation maps*: pinpoint the layers in a network that lead to the output

I would be very happy if every topic was chosen by at least one group. However, you are also welcome to explore a different topic of your choice.

Deliverables

You should submit the following in a PDF produced from LaTeX source code:

- Your name (and your partner's name if you want to work in a pair)
- A tentative title of your project

- The problem statement
- The data sets you plan to use
- Descriptions of the deep learning techniques you'll apply including 1) architecture, 2) loss, and 3) training method