EE 519: Deep Learning Theory & Fundamentals **Spring 2023**

Mini Project 3

Due: June 16, 2023, 11:59PM PT

Student Name: *Instructor Name: John Lipor*

Introduction

Diagnosing medical pathologies is a key application of artificial intelligence, especially in recent years as the success of convolutional neural networks (CNNs) has become increasingly prominent. One open research problem in this area is to determine pathologies from ultrasound images, which are sensitive to blood flow, as well as other tissues that would not be visible through other imaging modalities. As a step toward this goal, in this project, you will work with ultrasound eye images to predict the horizontal and vertical diameter of the eyeball. An example image is shown in Fig. 1, where the human-labeled diameters are also shown.

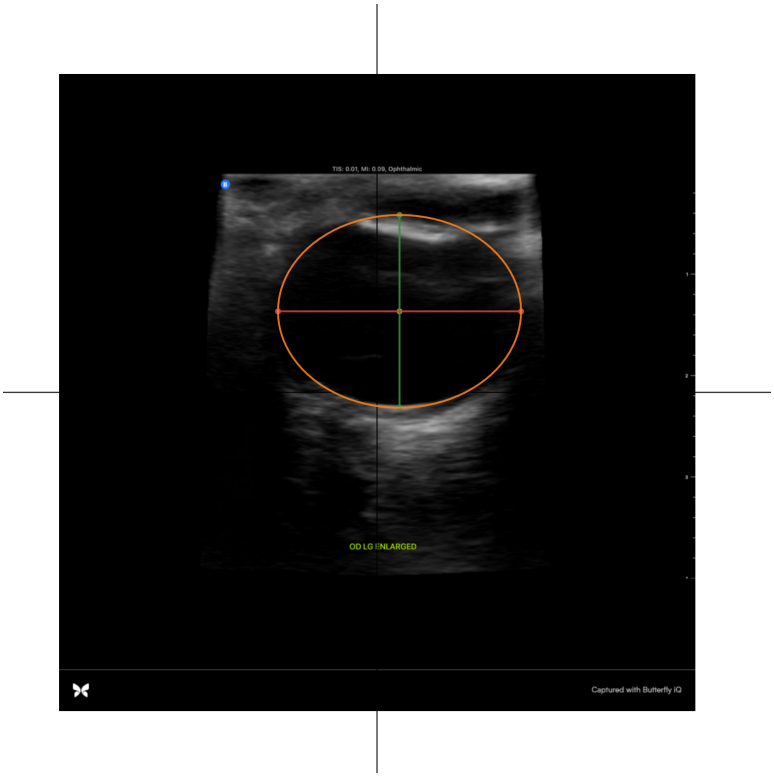

Figure 1: Example ultrasound image with horizontal and vertical diameters drawn. The goal of this project is to automatically determine these diameters from the raw image.

You are provided with ultrasound images and their corresponding labels. However, since we are attempting to predict both the horizontal and vertical diameters, this is a multi-output regression problem. Further, since labeling these images is a costly process, we are again working in a data-limited regime with only 100 total images, 80 used for training and 20 for testing. Your task is to train a neural network to predict both

diameters. As with other projects, you will submit your predictions on the test set, and I will evaluate your score using the mean absolute error. [1]

Notes

1. Your network must be coded and trained using PyTorch and D2L (not Keras or other similar packages).
2. You may make full use of StackExchange, Medium, or any other online resource you can find.
3. You may copy code from the internet, but you must **understand** what the code is doing and **cite** the source you copied it from.
4. You may **not** copy code from other students in this class. You **may** share MP1 or MP2 code among yourselves.

Requirements

You must create a report in either LaTeXor an approved type setting language (**not** Microsoft Word) that contains the following sections:

1. Problem description
   - Explain what data you have to work with, what algorithms you will use, what your goal is, and why anyone should care.
2. Exploratory data analysis (EDA)
   - Since we are working with images, EDA will look different and perhaps be less involved than for MP1. You should still examine the ranges of image values and labels, remove outliers, etc.
   - There are many excellent tutorials on EDA online. Cite whichever ones you found helpful.
   - Be sure to make any significant findings stand out and try to keep this concise.
3. Challenges
   - What were the challenges you encountered when applying machine learning to this dataset?
   - Did these challenges mainly result from the data? From results? Installing libraries to perform preprocessing?

[1] The top scoring project will receive 507 Lipor points.

4. Approach

   - **(NEW)** This is a multi-output regression problem, which you must solve using a single network. In other words, you may not train separate networks for vertical and horizontal diameters. Instead, you must structure your network to have two separate outputs. You may have as many (or few) shared layers as you wish.
   - Upon opening the dataset, you may notice that there are very few images to use for training, which will be further reduced when you create training and validation sets. As a result, you **must** describe approaches you took to overcoming the limited dataset size. This need not be through image augmentation,[2] but it could be.
   - Provide a thorough but concise description of your approach, including (but not limited to) your approach to wrangling, preprocessing, and improving the performance of your predictor.
   - To avoid overfitting, you must use some form of validation (either a single validation set or cross validation) to tune your hyperparameters. Describe how you went about this and why.
   - A summary paragraph at the beginning of this section could be helpful.

5. Evaluation and summary

   - You must provide some insights into which parts of the images your CNN is utilizing for prediction. For example, you may use class activation maps or guided backpropagation to visualize this.
   - You may also wish to display the learned feature maps and filters, but only do so if they are informative and you are able to draw a conclusion from them.
   - Make use of any meaningful metrics, regression MSE, classification error), but only include metrics that tell you something different.
   - Consider diving into the data to see if you can determine any trends regarding which points were misclassified.
   - Likewise consider if there were any features that were particularly important or unimportant.
   - Summarize your solution, describing what worked, what didn't, what the main limitations are (e.g., if nonlinear classifiers would improve results), and any general conclusions about the dataset.

6. What I learned

   - Describe the main skills/tools that you learned and used for this project and how you learned them.

[2] See `torchvision.transforms`.

Grading

The grading breakdown is below. Each section will be graded according to technical correctness, effort, and creativity. Note that clarity of writing is a major component. You should put yourself in the place of writing for a boss or senior in a workplace. If your writing is terrible, they will assume your work is terrible.

Item	**Percentage**
Description	5%
EDA	20%
Challenges	10%
Approach	30%
Evaluation	20%
What I learned	5%
Clarity/conciseness of written communication	10%