Deploying Deep Learning Models for Organ and Lesion Segmentation
Jed Perler, 2021

This summer, I had the great opportunity to intern in New York City at the Mount Sinai Hospital Department of Neurosurgery as a part of the Sinai BioDesign team. This group, directed by Bowdoin alumnus Anthony Costa PhD, works in conjunction with physicians and researchers to innovate new medical devices and identify and develop technological solutions to identified clinical needs. The BioDesign lab is equipped with 3D printers, milling machines, and other benchtop tools that provide efficient design and in-house prototyping capabilities to innovate new medical devices. During my time in the lab, I participated in the design and fabrication processes of several projects. For example, I used CAD programs to digitally model prototypes, a laser cutter to create the physical devices, and acrylic welding to finish prototypes.

In addition to these facilities, the BioDesign group also has a large computational research effort, called AISINAI, primarily in the area of artificial intelligence (AI) applied to medicine. This is where I dedicated most of my time in the lab. My primary project for the summer focused on addressing some of the issues that researchers encounter when applying AI to healthcare. One of the most important basic roles for AI in healthcare today is to leverage machine learning models to automatically segment organs, tumors, and other lesions from medical images such as MRIs or CT scans. Having these kinds of machine learning models that can quickly analyze medical images can benefit hospitals by allowing faster treatment in time-sensitive cases, as well as more accurate diagnoses. However, in order to create a model that can perform well enough to be deployed in the real world, an artificial neural network algorithm must be trained with a large set of medical images that have the organ or lesion of interest correctly labeled or segmented by a physician. Obtaining data sets large enough for effective training is often difficult, as there is often great time and cost associated with generating labeled medical images in the clinic. Even when a model is preforming well, it is often difficult to deploy it in a way that is easy for physicians to use. These issues inspired my work, which used Clara, a software development kit (SDK) developed by NVIDIA. Using this kit, we worked to construct an AI Assisted Annotation Server that allows for automatic segmentation and assisted annotation of medical images in real time. When one of these servers is constructed, a trained artificial neural network model can be pushed on to the server and physicians can easily interact with it in the clinic to produce automatic segmentations in seconds to aid their diagnosis and treatment. Additionally, the server allows physicians to use commonplace visualization software to create new labels on blank images in seconds, as opposed to a slice by slice manual labelling process that can take anywhere from ten to thirty minutes. These newly labeled images can then be used to create larger and more accurate training sets.

As part of my internship, I also had the opportunity to attend and participate in meetings, seminars, and research presentations that exposed me to the range of work going on in the field of AI in healthcare.

Faculty Mentor: Stephen Majercik

Funded by the Computer Science Student Fellowship Award