Predicting Sustainability Measures from Satellite Imagery using Deep Learning
Ben Bockmann, 25

This past summer I had the pleasure of participating in a research project led by Professor Farias and Professor Nascimento. The goal of the project was simple: to create a Deep Learning model that predicts a particular piece of demographic data for a given region. As a starting point, we endeavored to make a model that, given a satellite image of a zip code, predicts that zip code's poverty rate.

We began the summer with our focus directed towards the state of Maine. We thought its wide range of population densities between its zip codes would create a rich dataset. My research partner, Justas Bardauskas, used a Google Maps API to get images of all the zip codes in Maine. Using freely available data, he compiled the poverty rates for each county, and then associated each zip code with a county. Essentially, our final dataset consisted of aerial imagery at the zip code level paired with a poverty rate.

The first step in creating a Deep Learning model using a neural network is to find a suitable set of pretrained weights. Basically, a neural network consists of a set of weights that can be tuned for a given task. By taking advantage of existing weights, we were able to apply transfer learning to our task. Transfer learning drastically improves the accuracy of a model by harnessing large amounts of training that would not otherwise be feasible. We found a Python library called TorchGeo that contained many sets of pretrained weights, as well as modules that can be used for different types of Deep Learning tasks. We settled on a set of weights that had been trained on large amounts of RGB satellite imagery, which matched the type of images contained in our dataset. For the purposes of our model, we used the TorchGeo regression task, as we wanted the model's predictions to be fluid and not bound by existing poverty rates.

To program the model, I wrote two custom classes using Python. The first facilitated the loading of images from a Google Drive, reading in of files containing the poverty rates, and the pairing of images and poverty rates. The class split images into patches, which improved the accuracy of the model. We found that, for the state of Maine, splitting each zip code into 64 even patches resulted in the best poverty prediction accuracy. When we expanded our dataset to include the states Rhode Island and Oregon, we found that 49 patches per image produced the best result.

One area of the project that went unfinished was an attempt to filter out unhelpful images from the dataset. To do this, I wrote models that used deep learning to classify images into fields such as "Forest", "Water", or "Farm Land". Our idea was that by removing images that were water or forest, we could simultaneously reduce the cost and time of training and improve our models accuracy. Unfortunately, we were not able to put this feature into the final model.
Directions for further research

The next step for this project is to expand the data included in each zip code. Our idea to accomplish this involves the Google Places API, which can be used to get data about the businesses in a given area. The data for each business includes things such as price and business type. Including these additional data in the model would provide relevant information to the model, and we hope that it would further increase its accuracy. I hope to continue this research in an Independent Study at Bowdoin before I graduate.

Faculty Mentor: Jeová Farias
Funded by the Surdna Foundation Undergraduate Research Fellowship