

our final dataset consisted of aerial imagery at the zip code level paired with a poverty rate.

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

□

The first step in creating a Deep Learning model using a neural network is to find a suitable set of pretrained weights. Basically 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 t s to get data