

Poverty Prediction Deep Learning Model

Mingi Kang, Class of 2026

This summer, under the Christenfeld Summer Research Fellowship, I collaborated with Professor Jeova Farias on a project aimed at developing a deep learning model to enhance the accuracy of satellite imagery classification. Our focus was on a novel technique: *Convolutional Nearest Neighbor (ConvNN)*, which aims to improve image classification by analyzing the relationships between pixels and their nearest neighbors rather than relying solely on adjacent pixels. This approach has the potential to significantly improve traditional methods used in image recognition, especially in complex tasks such as poverty prediction in satellite images.

The primary objective of our project was to create a deep learning model that could be applied to satellite imagery to predict poverty levels. The ConvNN approach seeks to improve the accuracy of image classification by incorporating a broader context within the images, particularly by analyzing neighboring structures and landmarks. Our hypothesis was that this broader analysis would provide a more comprehensive understanding of the data, ultimately leading to better predictions.

To achieve this, we first implemented and tested the ConvNN model on simpler datasets before applying it to the more complex task of poverty prediction. We began with a 1D model (*Conv1d_NN*) using the MNIST1D dataset, a 1-dimensional representation of handwritten digits. Although the standalone Nearest Neighbor model did not outperform traditional convolutional networks, a hybrid model that combined both ConvNN and traditional convolutional techniques showed significant performance improvements. This hybrid model demonstrated that our approach could effectively leverage the relationships between central pixels and their nearest neighbors.

Following the success of the 1D model, we transitioned to a 2D implementation (*Conv2d_NN*), which included additional customizability options such as changeable channels, pixel shuffle/unshuffle capabilities, and varied sampling methods. We tested the 2D models on well-known datasets such as MNIST, Fashion MNIST, and CIFAR10, evaluating their performance on both classification and denoising tasks. While traditional convolutional models often outperformed our ConvNN variants, the results were encouraging and pointed to areas for further refinement.

This summer represents the first phase of a broader project aimed at developing a Poverty Prediction Deep Learning Model. Moving forward, we plan to continue refining our ConvNN models throughout the 2024-2025 academic year, with the goal of optimizing their performance for various image processing tasks. The aim is to apply the ConvNN technique to predict poverty levels using satellite imagery, with the second phase of the project expected to be completed by the end of the 2024-2025 school year.

In conclusion, our ConvNN project has laid a strong foundation for the exploration of new convolutional techniques in deep learning, demonstrating the potential for enhanced image classification accuracy. We are excited to continue this work and explore its application in real-world scenarios, particularly in assessing poverty levels in underdeveloped regions. Further details on our methods and code can be found in our GitHub repository linked below.

This research was made possible through the support of the Christenfeld Summer Research Fellowship, Bowdoin College, and the guidance of Professor Jeova Farias.

Faculty Mentor: Jeova Farias

Funded by the Christenfeld Summer Research Fellowship

References: <https://github.com/mkang817415/Convolutional-Nearest-Neighbor>