Deep Learning for Social Sciences - Assignment: Land Use Classification from Satellite Imagery

Giordano De Marzo

15 May 2025

To be handed in by 04.06.2025 23:59

Overview

In this assignment, you will use a Convolutional Neural Network (CNN) to classify land use patterns from satellite imagery. This task connects directly to applications including urban planning, environmental management, and the study of human-environment interactions. The EuroSAT dataset contains satellite images of different land use categories across Europe.

This assignment is closer to real-world applications of deep learning, so don't expect perfect accuracy. Some classes may be more challenging to distinguish than others due to visual similarities in the satellite imagery.

Complete tasks for receiving points, the maximum number of points you can get is 30. You can get bonus points completing the optional tasks. Discuss the bonus tasks in the report, you have up to 1.5 additional pages for the bonus tasks. Bonus points will not increase your mark over 30, but they can increase it up to 30 if you made some errors in the mandatory tasks.

Tasks

1. Data Import, Visualization and Preparation (10 points)

- Download the EuroSAT RGB dataset from `https://github.com/phelber/EuroSAT` (download link: `https://zenodo.org/records/7711810/files/EuroSAT_RGB.zip?download=1`)
- Preprocess the data and visualize some images with the corresponding labels to be sure everything is working
- Split the dataset into training (70%), validation (15%), and test (15%) sets

2. CNN Model Building, Training and Testing (10 points)

- Build a CNN to classify the images into the 10 land use categories
- An accuracy above 85% is considered a good result for this dataset
- Train and test the model using techniques for improving performance, such as:
  - Data augmentation
  - Early stopping
  - Regularization techniques (dropout, L1/L2 regularization)
  - Appropriate optimization algorithms

3. Report Writing (10 points)

- Write a report (max 3 pages) detailing the procedure you followed and the architecture you chose
- The report must contain the learning curves of your best model both on the training and validation set
- Clearly report the performances of the model on the test set (accuracy, F1 Score, classification report, confusion matrix)
- Identify which classes were most difficult to classify and hypothesize why. Report some examples of misclassifications
- A person must be able to replicate all your analysis by only reading your report

Bonus Tasks

Bonus Task 1: MLP vs CNN comparison (2 bonus points)

- Train a Multi-Layer Perceptron (MLP) to classify the images and compare its performance with the CNN

Bonus Task 2: Feature visualization and interpretation (2 bonus points)

- Visualize the activations of the last convolutional layer and discuss which features the model has learned
- Show some relevant examples in the report and interpret what patterns the CNN is detecting for different land use categories
- Discuss how these visualizations help understand the model's classification decisions