Assignment 9: Train CNN for Scene Labeling Task

40 Points

Deliverable: Notebook (.ipynb file) with all required code to complete the stated tasks. Answer all questions in Markdown cells.

Overview: The goals of this assignment are to train the CNN architecture defined in Assignment 8 using the data prepared in Assignment 7 and assess the model predictions using the withheld test data.

Data:

This assignment makes use of the *physioDL* dataset, which is available on FigShare:

Maxwell, A.E., 2024. physioDL: A dataset for geomorphic deep learning representing a scene classification task (predict physiographic region in which a hillshade occurs). https://doi.org/10.6084/m9.figshare.26363824.

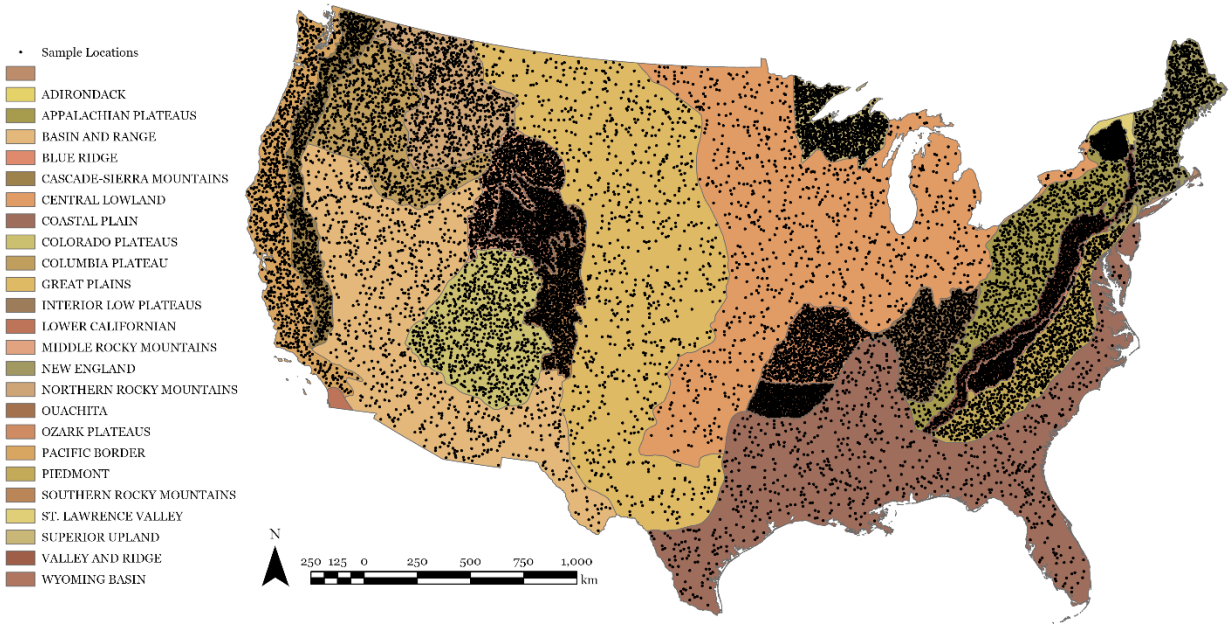

The task presented in this dataset is to predict the physiographic province of an area based on a hillshade image. Terrain data were derived from the 30 m (1 arc-second) 3DEP product across the entirety of CONUS. Each chip has a spatial resolution of 30 m and 256 rows and columns of pixels. As a result, each chip measures 7,680 meters-by-7,680 meters. Two datasets are provided. Chips in the *hs* folder represent a multidirectional hillshade while chips in the *ths* folder represent a tinted multidirectional hillshade. You will use the data in the *hs* folder. Data are represented in 8-bit (0 to 255 scale, integer values). Data are projected to the Web Mercator projection relative to the WGS84 datum. Data were split into training, test, and validation

partitions using stratified random sampling by physiographic province. 70% of the samples per region were selected for training, 15% for testing, and 15% for validation. There are a total of 16,325 chips. The following 22 physiographic regions are represented: "ADIRONDACK" , "APPALACHIAN PLATEAUS", "BASIN AND RANGE", "BLUE RIDGE", "CASCADE-SIERRA MOUNTAINS", "CENTRAL LOWLAND", "COASTAL PLAIN", "COLORADO PLATEAUS", "COLUMBIA PLATEAU", "GREAT PLAINS", "INTERIOR LOW PLATEAUS", "MIDDLE ROCKY MOUNTAINS", "NEW ENGLAND", "NORTHERN ROCKY MOUNTAINS", "OUACHITA", "OZARK PLATEAUS", "PACIFIC BORDER", and "PIEDMONT", "SOUTHERN ROCKY MOUNTAINS", "SUPERIOR UPLAND", "VALLEY AND RIDGE", and "WYOMING BASIN".

physioDL.csv: Table listing all image chips and associated physiographic province (id = unique ID for each chip; region = physiographic province; fnameHS = file name of associated chip in *hs* folder; fnameTHS = file name of associated chip in *ths* folder; set = data split (train, test, or validation).

chipCounts.csv: Number of chips in each data partition per physiographic province.

Tasks:

T1: Prepare your data as defined in Assignment 7. (5 Points)

T2: Define and instantiate your model as defined in Assignment 8. (5 Points)

T3: Train the model using the following configuration. (10 Points)

1. Use a mini-batch size of 32.
2. Train the model for 50 epochs.
3. Use the AdamW optimizer with a learning rate of 1e-3.
4. Use a cross entropy loss.
5. Monitor the loss for the training and validation sets at the end of each training epoch.
6. Monitor the overall accuracy and class aggregated macro-averaged F1-score for both the training and validation data.
7. Save the model that provides the best F1-score for the validation data.

T4: Assess the model using the test set. Calculate overall accuracy and class aggregated macro-averaged recall, precision, and F1-score. Also, calculate the confusion matrix and all class-level precision, recall, and F1-scores. (10 Points)

T5: Discuss the modeling results including: (10 Points)

1. Learning process stability
2. Evidence of overfitting
3. Overall model performance on the training and validation sets
4. Overall model performance on the test set
5. Class-level performance on the test set

6. What physiographic regions were most confused or difficult to distinguish using the single band hillshade image?