

Exercise 5: Image classification - Part 2

Introduction

In this exercise, we will perform image classification using regions/superpixels as spatial support, and compute the accuracy of the predictions. Please use the Jupyter Notebook provided to complete this exercise.

Info

All parts in the code that require your input for completion are marked with flag “#TODO”.

Tasks

1 Setup

Open the file `ex5.ipynb` (available on Moodle) using Jupyter.

2 Extract regions and their features

- 2.1 Download the dataset
- 2.2 For each image in the dataset, compute regions and their features using the function implemented in the previous exercise

3 Create a training dataset

- 3.1 Read the images from the training set and compute the ideal label for each superpixel (the label that has the largest intersection with each region)
- 3.2 Create an array of targets (ground truth), per region, that joins the targets of all the training images
- 3.3 Create an array of features that joins the features of all the training images

4 Normalize features

- 4.1 Normalize features by subtracting the mean and dividing by the standard deviation

5 Train a Random Forest classifier with default parameters

- 5.1 Create a Random Forest classifier, using the library `sklearn`
- 5.2 Use the function `fit` to train the classifier with the normalized features and targets

6 Predict classification map

- 6.1 Predict the classification map with the current trained classifier (with default parameters)
- 6.2 Compute the accuracy and confusion matrix on the validation set

7 Compute accuracy and the confusion matrix

- 7.1 On the images of the validation set, compute the accuracy and the confusion matrix

8 Search for good parameter values in the validation set

- 8.1 Compute the accuracy in the validation set for several models trained with different values of `n_estimators` and `max_depth`
- 8.2 Select the parameters that obtain the best accuracy in the validation set

9 Predict classification maps with the best parameters found in the validation set

- 9.1 Train a Random Forest model with the parameters that performed best in the validation set
- 9.2 Compute the accuracy and confusion matrix on the test set

Q Observing the confusion matrix. Which types of errors do you find in the predictions?

10 Visualize predictions

- 10.1 Visualize one of the classification predictions saved on disk
- 10.2 Visualize the ground truth label map that corresponds to the classification prediction that you selected to plot