

Exercise 4: Image classification - Part 1

Introduction

In this exercise, we will perform image classification using regions/superpixels as spatial support. 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 `ex4.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 fuction 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 (groud 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 substracting the mean and dividing by the standard deviation

5 Train Random Forest classifier

- 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

For each image in the test set

- 6.1 Load features
- 6.2 Predict labels by region
- 6.3 Predict labels for each pixel
- 6.4 Save the predictions on disk

7 Visualize predictions

- 7.1 Visualize one of the classification predictions saved on disk, using the function `display_label_image` defined in the Jupyter Notebook
- 7.2 Visualize the ground truth label map that corresponds to the classification prediction that you selected to plot

8 Answer the following question

Q Comparing the prediction and the ground truth, which types of errors do you find in the predictions?