

Computer Vision

CSCI-GA.2272-001

Assignment 3

November 7, 2016

Overview

This assignment requires you to participate in a Kaggle competition with the rest of the class on the The German Traffic Sign Recognition Benchmark [<http://benchmark.ini.rub.de/?section=gtsrb&subsection=news>]. The objective is to produce a model that gives the highest possible accuracy on the test portion of this dataset. The Kaggle competition is invite only and you will receive an invitation shortly.

You should first clone the Github repository <https://github.com/rajegannathan/traffic-sign-detection>. The `main.lua` file contains code for training and evaluating your models. Once training is completed, it produces a file `submission.csv` that lists the IDs of the test set images, along with their predicted label. This file should be uploaded to the Kaggle webpage, which will then produce a test accuracy score.

The model architecture is specified in `model.lua`. Currently a simple baseline model is provided. You can use this model or create your own. Make sure to use the appropriate `-model` command line argument to select the correct model for training and testing. You are free to implement any approach covered in class, or in the research literature. Of course, tricks that you devise yourself are also encouraged.

The test and training data are provided in the Kaggle competition data download page. You should unzip the images into `data/train_images` and `data/test_images` subdirectory which is already provided in the code.

Requirements

You should perform this assignment in Torch, modifying the template code provided.

This assignment is due on **Thursday December 8th** at 11.59pm. Please note that the late policy is as follows: (a) assignments that are late by less than 24hrs will suffer a 10% reduction; (b) those between 24 and 72 hrs late will suffer a 25% reduction and (c) those more 72hrs late will suffer a 50% reduction. You are strongly encouraged to start the assignment early and don't be afraid to ask for help from either the TA or myself.

You are allowed to collaborate with other students in terms discussing ideas and possible solutions. However you code up the solution yourself, i.e. you must write your own code. Copying your friends code and just changing all the names of the variables is not allowed! You are not allowed to use solutions from similar assignments in courses from other institutions, or those found elsewhere on the web.

Your solutions should be emailed to me at (fergus@cs.nyu.edu) in a single zip file with the filename: `lastname_firstname_a3.zip`. This should include: (i) a brief report explaining the model architectures you explored, citing any relevant papers or techniques that you used. You should also include convergence plots of training accuracy vs epoch for relevant models. (ii) the model file for your best submission. Please run `clearstate` before saving it as it will drastically reduce the file size. If the model is too big to be emailed, please upload it to CIMS or Google Drive or similar and provide a link to it in your report.

Important Details

- You are only allowed eight (8) submissions to the Kaggle evaluation server per day. This is to prevent over-fitting on the test dataset. So be sure to start the assignment early!
- You are **NOT ALLOWED** to use the test set labels during training in any way. Doing so will be regarded as cheating and penalized accordingly.
- The evaluation metric is accuracy, i.e. the fraction of test set examples where the predicted label agrees with the ground truth label.

- You should be able to achieve a test accuracy of at least 0.95. However, top grades will be reserved for entries that exceed this by a significant margin.