Homework Assignment #7

DUE: 5:00pm, Sunday May 15th (Electronic turnin required)

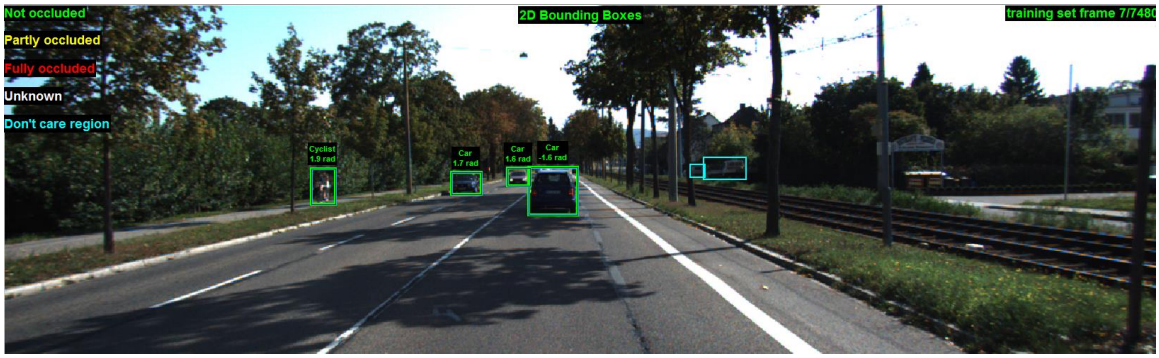

You are to use some public-domain datasets to perform certain object recognition tasks. For a homework assignment, you might want to look into small data sets, e.g., http://www.vision.caltech.edu/Image_Datasets/Caltech101/. This data set contains pictures of objects belonging to 101 categories. About 40 to 800 images per category. Most categories have about 50 images. The size of each image is roughly 300 x 200 pixels.

Larger and more meaningful datasets are also available. For example, the KITTI's object detection and object orientation estimation benchmark (http://www.cvlibs.net/datasets/kitti/eval_object.php) consists of 7481 training images and 7518 test images, comprising a total of 80,256 labeled objects.

This is a wide-open assignment, more like a mini project, as we are not dictating (1) what data sets you choose, (2) what features you use, (3) how you partition your data sets for training and validation, (4) what classification scheme you use, (5) how you perform cross validation, etc. You need to write a short report (< 5 pages) that explains your choices of (1) to (5) above. In your report, you also need to detail your experimental results, e.g., precision vs. recall tradeoff (ROC curve), expected recognition accuracy, and lessons learned.

Hint: There are many components in a visual recognition experiment. Some essential steps are

1. Select and collect training data
2. Design and implement feature extractors
3. Design and implement a classification scheme
4. Training: determine parameters of the classifier and training error. Usually, this requires labeled data with some cross-validation schemes (n-fold)
5. Validation: determine generalization error (i.e., error on labeled data that are not used in training)
6. Deploy the system for real-world use

This homework is for you to "complete the loop" (at least steps 1-5 above) and learn how such an experiment is conducted.

It is recommended that you try to exploit available resources as much as possible. For example, you may decide to do an experiment on texture image classification. In that case, try to find a data set online (e.g., http://homepages.inf.ed.ac.uk/rbf/CVonline/Imagedbase.htm), and use an existing feature scheme and its public-domain implementation (e.g., Garbor wavelets). Then select a classifier (e.g., implement a simple linear discriminant classifier or use an existing SVM package). Perform n-fold validation experiments and record the training error (if possible, try varying some parameter setting to generate an ROC curve). Finally, validate your feature/classifier design on unseen, labelled data as a reality check to see how your scheme generalizes in real world.

In the above scenario, you did not do a lot of implementation or nothing at all. But you will go through the whole loop in such an experiment in order to generate a report. Finally, this is a "getting-your-feet-wet" experimentation. Hence, It is not necessary that you obtain an excellent recognition rate. In fact, it is well known that simple classifiers such as linear discriminant function may not achieve good results, so do not "doctor" data to report unrealistic recognition rate.