Laboratory Experience 2 - Mandatory

Task Goal of this experience is to get familiar with the kernel combination techniques most commonly used for object and scene recognition.
You are asked to replicate the results in [3], in particular to perform experiments using two object recognition datasets:

Oxford Flower dataset [4]. A flower recognition task, consisting of 17 species of flowers, with 80 images of each

Caltech 101 dataset [2]. A 101 categories object recognition task (plus one category for background), generated by retrieving relevant images from the Google Image Search engine

and four different learning algorithms:

SVM with the χ^2 kernel computed on each single image descriptor

SVM with the average kernel [3] using all the features

SVM with the product kernel [3] using all the features

Multiple Kernel Learning [1] using all the features

The results obtained with techniques other than these don't need to be replicated.

Code

1. Download Shogun toolbox (v2.0.0), an open source library for machine learning:
   http://www.shogun-toolbox.org
   You need to compile the library (with the Matlab interface); from the prompt execute:

```
>> cd /path/to/shogun2/src/
>> ./configure
>> make
```

   Moreover, in order to make Matlab link to the correct libraries, before starting it (on bash) perform a:

```
>> export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/path/to/shogun2/src/shogun/:\
/path/to/shogun2/src/shogun/ui/
```

2. Download libsvm, a free implementation of Support Vector Machines:
   http://www.csie.ntu.edu.tw/∼cjlin/libsvm/
   You need to compile the Matlab interface; from the Matlab prompt execute:

```
>> cd /path/to/libsvm-3.12/matlab/
>> make
```

3. Download the MklExperience Matlab class and auxiliary files:
   http://www.idiap.ch/ftp/courses/EE-700/material/experience2/MklExperience.tar.gz
   Open MklExperience.m and modify the paths of the libsvm-mat and shogun libraries:

```
>> LIBSVMPATH='/path/to/libsvm-mat';  %path to the lib-svm solver matlab interface
>> SHOGUNPATH='/path/to/shogun2/src/interfaces/matlab_static'; %path to the shogun
```

   Inside the class you will also find an example of usage.

4. Complete the code where necessary.

Oxford Flowers

1. Download the Oxford Flowers dataset:
   http://www.robots.ox.ac.uk/∼vgg/data/flowers/17/

   - You need the data splits, the χ^2 distances CVPR 2006, and the χ^2 distances ICVGIP 2008.

2. Repeat the object recognition experiments described in Table 2 of [3], using C=10 for all methods. Run each experiment on the three splits and report the results as (mean $\pm$ std).

3. To measure the multiclass accuracy use the mean class recognition rate, over all classes and splits

4. For all the methods report also the average training/testing time.

Caltech 101

1. Download the Caltech 101 dataset:
   http://www.idiap.ch/ftp/courses/EE-700/material/experience2/Caltech101.zip
   It includes four directories:

   - n5_N1, n15_N1, n30_N1: each of these directories contains the precomputed training and testing kernel matrices for 7 features (V1S+, PHOG, Region covariance, LBP, SIFT, PHOW, Self-similarity), with a different number of training samples per class: 5, 15 and 30.

   - labels: training and testing labels

2. Consider 5, 15, and 30 samples per class and for each of them report the results in a table similar to Table 2 of [3]. Again, use C=10 for all methods.
   We consider only a single split of the data (N1: the original first split), so in this case you will have indicative results without standard deviation.

3. Changing the number of training samples, plot also the curves of the recognition rate obtained by the best feature, the average and product kernels and MKL, as in Figure 2(d) of [3].

4. For all the methods report also the training/testing time.

Comment all the results w.r.t. the computational costs and the performances achieved.

References

[1] F.R. Bach, G.R.G. Lanckriet, and M.I. Jordan. Multiple kernel learning, conic duality, and the smo algorithm. In *Proceedings of the twenty-first international conference on Machine learning*, page 6. ACM, 2004.

[2] L. Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, 106(1):59–70, 2007.

[3] P. Gehler and S. Nowozin. On feature combination for multiclass object classification. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 221–228. IEEE, 2009.

[4] M-E. Nilsback and A. Zisserman. A visual vocabulary for flower classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1447–1454, 2006.