

381V Visual Recognition, Spring 2016
Assignment 2: Recognizing object categories
<http://vision.cs.utexas.edu/381V-spring2016/>

Out: Monday Feb 22
Due: Wednesday Mar 9, 11:59 pm,
with follow-up result due Thursday Mar 10, 11:59 pm (see below)



Goal

For this assignment, the goal is to implement an image recognition system based on methods we have studied in class so far. You will evaluate your system on natural images containing 25 object categories from ImageNet. Feel free to use existing code for *components* of your system; however, it is not permissible to use existing executables or complete end-to-end systems you might find online to compute your results. Cite any components you have borrowed clearly in your writeup.

There will be extra credit for the top 3 best performing systems, as gauged on a held-out test set we will provide later (see below).

Data

We have a local copy of the ImageNet challenge data here:

```
/projects/cs381V.grauman/
```

which you can read directly for loading images for this project. There is also some limited additional temporary disk space available there for writing features, if you need it.

We will be operating with a subset of ImageNet for this assignment, consisting of 25 categories. So that results are comparable, you are to use the subset of training and test images specified in this file: `/projects/cs381V.grauman/ilenames.mat`. This mat file contains these variables:

- `classnames` is a cell array of 25 strings. It enumerates the category names for the 25 categories we randomly selected for inclusion in this assignment.
- `trainImNames` is a 25 x 100 cell array of JPEG file names. The $\{i,j\}$ -th entry in `trainImNames` specifies the j -th training image for the i -th selected category. There are 100 training images for each of the 25 categories. Your code will need to read this list to load the training images, which as the file names show, are located in subdirectories of `/projects/cs381V.grauman/`. Note that `classnames{i}` gives the descriptive label for all images specified by `trainImNames{i,:}`.
- `test1ImNames` is a 25 x 20 cell array of JPEG file names, formatted as above. The $\{i,j\}$ -th entry specifies the j -th testing image for the i -th category.
- `extraTrainImNames` is a cell array formatted like `trainImNames`, and points to extra training images that are optional for use.

What to implement

You are to implement an image recognition system that trains models for the 25 provided categories, and can predict the class label for a novel test image. The approach you take must be derived from one of the papers and methods we have studied in class so far, but the details are up to you. We will evaluate the assignment based on 1) your design is technically correct, 2) it follows a basic paradigm we've studied in class, and 3) the accompanying report demonstrates careful analysis and understanding of the method you have employed. For example, your system might be based on one of the following:

- Bag of words representation, computed on sparse or dense local features, combined with a chosen classifier (e.g., SVM with a spatial pyramid kernel).
- Boosting with rectangular features computed from integral images (following Viola-Jones)
- GIST or Histogram of Oriented Gradients (HoG) descriptors and SVMs
- Pre-trained AlexNet CNN features, combined with a chosen classifier.
- Fine-tuned CNN descriptors
- Nearest neighbors (if you choose this one, we will expect to see some effort put on the feature design and/or distance function used for finding neighbors)

See the course website for links to available code, feature extraction toolkits, etc.

While your system must make use of the provided training images and report results on all the provided test images, you may optionally use the training images listed in `extraTrainImNames`.

What to report and how to submit

Train your system on the specified training images, and test it on the specified test images. Write a pdf report **called `report.pdf`** that includes the following:

- A description of precisely how your approach was implemented, listing the steps and details of any choices you made.
- Compute and display a 25 x 25 confusion matrix, and explain the result. Report the average of the diagonal of this matrix as a summary of the accuracy for the 25-way multi-class task. We will create a leaderboard using this average accuracy number.
- Discuss the overall performance of your system, such as where it works and where it doesn't, and briefly why. The `classnames` variable will be useful for qualitative analysis, and you may also want to examine image examples to get a feel for what is working or not.
- Test two variants of your system, and design a comparative experiment to show how they differ. Carefully *explain* the impact on accuracy. Illustrating the impact on accuracy could go beyond the overall accuracy number, such as impact on particular classes, types of examples, etc. *Note: please avoid simply reporting the impact of changing some hyper-parameter, unless there is an interesting conceptual explanation for the impact.*

Grades will be determined based on implementation completeness as well as the insights demonstrated by the report discussion. Use figures or images to make your point clear when appropriate.

Submit your tarred code and the `report.pdf` file together in a single zip file on Canvas. Please do NOT include any train/test data in the submission.

Withheld test set: due Thursday March 10

After the assignment deadline, on Thursday Mar 10, we will release a second test set on Piazza. You will need to run your approach, exactly as it was submitted (**no additional training, parameter tweaking, etc. allowed!**) and report the results via email to Kai, in the form of the average accuracy per class. Please put CS 381V in the subject line of the email. Those results are due by 11:59 PM Thursday March 10. We will base the contest/extra credit winners on these results.