CSI 5325 Assignment 5

Greg Hamerly

assigned April 1, 2008 (really!); due May 2, 2008

1 Expectations

For this assignment, you should work in a group of 3-4 people. Each group should plan together to distribute the work and turn in one report. Some people may do different parts of the assignment, but all should contribute equally.

This will be the last assignment for the semester, so it takes the place of the last three assignments, and carries three times the weight of previous assignments.

The writeups you do for this course should be of high quality; complete but concise. Please structure your writeups well, introducing ideas logically (rather than simply chronologically). Your writeup should be formatted as if you would submit it to a conference or journal. It should include the following sections in the given order:

- Introduction – Describe the problem(s) you are working on at a high level, and give the reader a summary of the results you found.
- Methodology – Explain your data sets, data preprocessing, the algorithms and techniques you apply, the hypotheses you are testing, and how you will evaluate the success of your experiments. Use tables and figures as are helpful.
- Results – Explain the results of the experiments, using graphs and tables and text. Then, analyze your results.
- Conclusion – Summarize the work again, and if applicable, describe further work that you think would build on what you've done.

1.1 Submitting your work

You should turn in all your work in printed format in class, as well as by email to `hamerly@cs.baylor.edu`. Your email should have attached a zip file which has a single folder that contains all your materials. Please DO include all source code with your emails, but DO NOT print out your source code to hand in (just print and turn in writeup).

1.2 Tools for the course

As suggested in assignment 1, your work should be composed in LaTeX, and I suggest looking at MATLAB for programming.

2 Natural image analysis

For this project, you should use machine learning to classify natural images. We hope to be able to observe the parking lot behind the Rogers ECS building to identify the number of parking spots available (0 - n).

For this, I suggest your group choose a machine learning method that you think is suited to the task. There will be several milestones you should accomplish:

- gather images – make sure this training/test set contains images under different conditions (morning/afternoon/night, cloudy/sunny, etc.)
- label images – choose a representation for the target
- choose a learning algorithm – I recommend trying more than one, and comparing results
- pre-process the data
- train the algorithms
- test on held-out data
- provide an internet- or phone-based service for your application

You should beware of having images that are too similar to each other in the training/test sets. For example, if you take an image every 5 minutes, and nothing changes between two images, and one image is in the training set and one is in the test set, then a 1-nearest neighbor classifier will have perfect prediction on that image, but it is not really an independent test.

3 Service delivery

You should at the very least create some web page that reports the latest information on the number of parking spots.

For a more interesting project, put together a method of sending text message information on request. Text messages can usually be sent by emailing to the phone number at the service provider.

An even cooler application would be using the voice XML service at Tellme.com. You can think of this like a web browser in your telephone which goes to your server, reads the VXML description, reads it to you, and allows you to make a menu selection (with voice recognition!). Based on your choice, it will fetch another VXML page, etc. See `http://studio.tellme.com`.

4 Privacy issues

If we are going to use images we capture on campus, we will be careful that:

- the images should only be viewed by people in this course for purposes of training and testing the system
- the images should not be made available to anyone outside the course
- the images and any residual data computed from them should be erased after the semester is over

Currently, we are waiting for permission to capture the images we're interested in.

5 Evaluation criteria

The following sections will be used for grading your assignment (these adapted from Charles Elkan's grading criteria):

Category	Points
Introduction insightful on background and motivation	3
Precise and reproducible description of technical work	3
Sensible implementation decisions (code reuse, choice of PL, etc.)	2
Sufficiently conclusive experiments (large datasets, lesion studies, etc.)	3
Well-designed, well-described, reproducible experiments	3
Understandable presentation of results, preferably graphical	3
Clear, correct analysis of results (including statistical significance)	3
Insightful discussion of all major experimental results	3
Exciting and useful results with general applicability	3
Appropriate organization (logical, not chronological) and easy-to-read plain writing style	3
Correct spelling, grammar, and choice of words	3
Total	32

A lesion study means taking parts out of the system you are using to identify what will cause a system's performance to degrade. For this assignment, you need not perform statistical significance tests, but careful analysis is still important.