CS-498 Homework 1

Due: October 6, 2012

You may do this homework in groups of up to three.

Place a 20x20 grid of squares, each representing 100m on edge, over a satellite image of the suburb of Ekali, in Athens, Greece. Careful use of the web to collect data is part of the exercise (for example, you could look at map websites to obtain images of Ekali). We do not require that the grid is exactly centered.

1) Build a swimming pool detector. You may use VLfeat or other methods to build image features, and you may use classifier training software (VLfeat has a classifier training method; you could also use SVMlight or LIBSVM). To do this, you will need to create positive and negative example training images.
    1) Download satellite images of an appropriate scale from the web
    2) Mark some example swimming pools by hand, drawing an axis aligned box of fixed size around the pool. It is a good idea to have a reasonable number of pixels (at least several hundred) on the pool. If you have too many pixels on the pool, your detector will be slow. If you have too few, it will be inaccurate. Groups may share examples.
    3) Check whether pools are of fixed or different sizes; can you search for just one size of pool, or do you need to search for more.
    4) Build features for your example boxes. We suggest a color histogram, or at least a measure of the percentage of blue pixels. It is easy to use VLFeat to compute HOG features.
    5) Train a classifier to tell positive example boxes from negative example boxes. You can do this using VLFeat, SVMLight or LIBSVM, or some other piece of software; VLFeat and LIBSVM can be called from MATLAB.
    6) Now build the detector, by walking over the image, computing features for each box, then presenting the box to the classifier. Use non-maximum suppression, as described in class. It is worth trying different step-sizes to see which performs better.
2) How accurate is your detector? you should check this by computing its recall and precision for various settings of detection thresholds. To do so, you will need to evaluate its detections by hand for some squares (do at least 5, but don't do all; it's a lot of work, and there's no point).
3) For each square, automatically make an estimate of the number of swimming pools in that square, using your swimming pool detector.