Pattern Recognition and Machine Learning

James L. Crowley

ENSIMAG 3 - MMIS Fall Semester
Lab Project 1: 12 October 2016

The objective of this project is to evaluate the effectiveness of detection of skin pixels using color histograms as a detector for human faces. Test data and ground truth are provided by the "FDDB: Face Detection Data Set and Benchmark Home" of the University of Massachusetts. The data set can be found at http://vis-www.cs.umass.edu/fddb/ and is described in the paper [Jain and Learned-Miller 2010] available for download from the course web site.

Skin pixels will be detected using a ratio of color histograms calculated from a subset of the benchmark data set. Evaluation will be performed using ROC curves that plot True Positive Rate vs False Positive Rate. You should compare the ROC curves for face pixels using both 3D RGB histograms and 2D normalized chrominance histograms computed from different subsets (folds) of the test data.

Ground truth for face detection is provided in the form of a list of ellipses that have been manually fit to the faces in each image. Most of the images contain skin regions that are not part of a face, as illustrated in the following image. This is will be a source of false positive detections in your evaluation. Face ellipses also contain non-skin regions such as hair. This will be a source of false negatives.

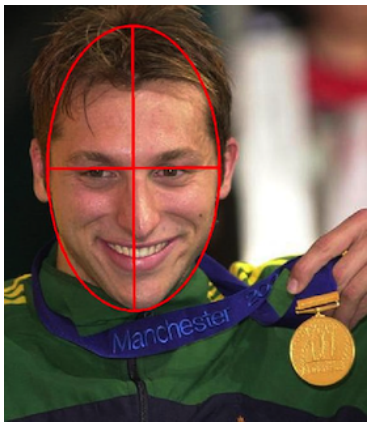

Each programming team should

1) Train a set of skin pixel detectors using both 2D and 3D histograms from sets of folds from the test data.
2) Plot ROC curves for the detectors using folds that were not used in training
3) Interpret the results, describing the effectiveness of the detectors and explaining the sources of errors.

Lab work will be reported with a written report in either French or English. Work will be evaluated based on the effectiveness of the experimental evaluations, and the clarity and depth of the explanation of experimental results.

[Jain and Learned-Miller 2010] V. Jain and E. Learned-Miller, "FDDB: A Benchmark for Face Detection in Unconstrained Settings", UMass Amherst Technical Report (2010).