

Pattern Recognition and Machine Learning

James L. Crowley and Nachwa Aboubakr

ENSIMAG 3 - MMIS
Lab Project 1:

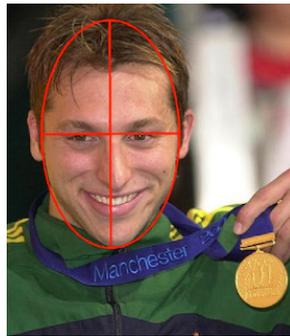
Fall Semester
10 October 2018

Face Detection with Color Histograms

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 “FDDDB: 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. Faces will be detected by summing the weighted skin pixel probabilities within a rectangular region of interest (ROI). ROIs are defined from the hypothesis that a face can be found at a position and size. Project teams should determine a range of size and positions from inspection of the test data. Evaluation will be performed using ROC curves that plot True Positive Rate vs False Positive Rate. Project teams 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. Pixels from within the ellipses can be used as training data for skin color histograms.

Most of the images contain skin regions that are not part of a face, as illustrated in the following image. This 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. Project teams should compare the effectiveness of different color codings for skin detection as well as different methods for face detection within a ROI.



Each programming team should

- 1) Train a detector for skin pixels using subsets (folds) from the test data.
- 2) Construct a sliding window face detector that sum probabilities in a ROI and decides Face/No Face for each position and size.
- 3) Plot ROC curves for the detectors using folds that were not used in training
- 4) Interpret the results, describing the effectiveness of the detectors and explaining the sources of errors.

A detected face is a True face if the center of the ROI is within the face ellipse given in the ground truth. It is possible to improve results by clustering adjacent detections.

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. Written Reports are due on 15 November.

Grading Scale for Lab project 1

The following is an indicative Barometer for Grading. Actual grades will depend on a subjective appreciation for the amount of effort deployed and the depth of understanding.

Grade	Example of Criteria
10	ROC plot showing results of face detection with RGB color histograms. Detection using sum of probabilities in ROIs. Tests with only a single training strategy. Reasonably clear description of experiments.
12	ROC plots comparing face detection with RGB and Chrominance (rg) histograms trained under different training strategies. Detection using un-weighted sum of probabilities in ROI. Tests with only a single training strategy. Clear description of experiments. Discussion of which technique worked better and why.
14	ROC plots comparing face detection with RGB and Chrominance (rg) histograms with different color quantizations, trained under different training strategies. Detection using weighted and un-weighted sum of probabilities in the ROI. Tests with multiple training strategies. Clear description of experiments. Analysis and explanation of common sources of errors.
16	ROC plots comparing face detection with RGB and Chrominance (rg) histograms with different color quantizations, trained under different training strategies. Detection using weighted and un-weighted sum of probabilities in ROI. Experiments with clustering of adjacent detections. Tests with multiple training strategies. Explanation if/why certain training strategies, color spaces work and histogram quantization work better. Insightful explanation of results.
18	ROC plots comparing face detection with RGB and Chrominance (rg) histograms with different color quantizations, trained under different training strategies. Detection using weighted and un-weighted sum of probabilities in ROI. Experiments with clustering of adjacent detections. Tests with multiple training strategies. Explanation if/why certain training strategies, color spaces work and histogram quantizations work better. Insightful explanation of results. Analysis of failures in face detection. Suggestions for improvements.
20	All of the above plus additional unexpected insights or results.

Creativity and originality will be rewarded.

[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).