

Pattern Recognition and Machine Learning

James L. Crowley

ENSIMAG 3 - MMIS
Lab Project 3:

Fall Semester
8 January 2020

Face Detection with Artificial Neural Networks

The objective of this project is to evaluate the effectiveness of different Neural Network techniques for building a face detector. The written report for this exercise is due on 5 February 2020.

Each programming team should

- 1) Implement and evaluate one or more neural networks for detecting faces.
- 2) Plot error rates for each detector when applied with variations in parameters using the FDDB data set as well as other benchmark data sets
- 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.

Programming teams are given considerable freedom in their choice of techniques to evaluate. You may use Neural Network code downloaded from the internet as well as your own implementations of code. It is essential to explain the origin and principles of operations for each evaluated face detector. 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 displayed in the results. Creativity is encouraged and will be rewarded!

Grade	Example of Criteria
10-12	Evaluation of a sliding window face detector using a 3-layer network trained with back propagation. ROC and/or precision-recall plots for the face detector using the FDDB data set. Explanation for learning and detection processes. Test with variations of parameters. Description of experiments. Discussion of results.
12-14	Implementation of a sliding window face detector using a 3-layer network with back propagation. ROC and/or precision-recall plots for the FDDB data set as well as other data sets. Explanation for learning and detection processes. Test with a range of parameters, window sizes (from 16x16 to 32x32) and training data sets. Clear description of experiments. Insightful discussion of results.
14-16	Comparison of a sliding window face detector using a 3-layer network to a multi-layer network. ROC and/or precision-recall plots for the FDDB data set as well as other data sets. Explanation for learning and detection processes. Comparison of results using different size windows (from 16x16 to 32x32) and using both gray-scale and RGB images as well as other parameters.
16-18	Comparison of sliding window face detectors using a 3-layer network, one or more multi-layer networks and a convolutional network. Comparison of results using different size windows (from 16x16 to 32x32) and using both gray-scale and RGB images as well as other parameters such as size of receptive fields for Convolutional net. ROC plots and other metrics comparing sizes of sliding windows, size and number of receptive fields, use of color and other variations of techniques and parameters.
18-20	All of the above plus additional unexpected insights or results.