

LFW Results Using a Combined Nowak Plus MERL Recognizer

Gary B. Huang¹, Michael J. Jones², and Erik Learned-Miller¹

¹ University of Massachusetts Amherst, Amherst, MA
{gbhuang, elm}@cs.umass.edu

² Mitsubishi Electric Research Laboratory, Cambridge, MA
mjones@merl.com

MERL's recognizer has two major components (Figure 1). The first is an aligner whose input is any arbitrary image and whose output is a cropped and rectified face if the image contains a face. The second is a comparator whose input is two aligned faces and whose output is a similarity score. The comparator is described in detail in [3].

The aligner first detects any faces in the image using a Viola-Jones frontal face detector. For any faces found, it tries to find nine facial landmark points by running nine Viola-Jones type facial landmark detectors. The facial landmark detectors are trained using the same Viola-Jones AdaBoost cascade method as the face detector. For each facial landmark, a set of aligned positive examples are manually marked. Some positive examples are shown for the right outside eye corner in Figure 2. The negative examples consist of other patches of faces and non-face patches.

After the facial landmark detectors are run, if at least six landmarks are found, then the optimal scale, in-plane rotation and translation in x and y are found which bring the landmark points into alignment with the landmark points of an average face. A failure to acquire error occurs when either a face is not detected or fewer than six facial landmark points are found. The similarity transform allows the input face to be cropped and rectified to a fixed size image (we use 64x80). After rectifying, some simple lighting normalization is applied. The normalization we use is similar to subtracting out the best fitting brightness plane ([4]).

Previous state of the art performance on the image-restricted unseen pair matching task of LFW was achieved by running the recognition system of Nowak and Jurie [1] using the aligned images produced by the funneling algorithm of Huang et al. [2]. By averaging this score with the score produced by the MERL system (after normalizing the scores to be on the same scale), we obtain new state of the art results, as shown in Figure 3. The two systems had a correlation coefficient of 0.5886. When the MERL aligner gave a failure to acquire error (occurring on about 11% of the LFW images), the funneled image was used, after applying a pre-defined transform to roughly match the images produced by the MERL aligner.




Fig. 1. Recognition steps of the MERL recognizer.



Fig. 2. Positive examples for right outside eye corner landmark.

References

1. Nowak, E., Jurie, F.: Learning visual similarity measures for comparing never seen objects. In: CVPR. (2007)
2. Huang, G., Jain, V., Learned-Miller, E.: Unsupervised joint alignment of complex images. In: ICCV. (2007)
3. Jones, M., Viola, P.: Face recognition using boosted local features (2003)
4. Rowley, H., Baluja, S., Kanade, T.: Neural network-based face detection. In: PAMI. Volume 20. (1998) 22–38

Method	$\hat{\mu} \pm S_E$
Nowak	0.7393 ± 0.0049
Merl	0.7052 ± 0.0060
Combined	0.7618 ± 0.0058




Fig. 3. Accuracy (left) and ROC curves on LFW.