

Aurora Face Recognition Technical Report: Evaluation of Algorithm “Aurora-c-2014-1” on Labeled Faces in the Wild

Dr T. Heseltine Dr P. Szeptycki Mr J. Gomes Dr M.C. Ruiz Dr P. Li

t.heseltine@auroracs.co.uk
Aurora Computer Services Ltd.
www.facerec.com

Abstract

We evaluate the performance of a new method of face recognition on the well-known benchmark Labeled Faces in the Wild (LFW) dataset [1]. The method, developed by the Core Technology Research Team at Aurora, achieves a mean classification accuracy of 93.24% on the unrestricted view 2 test set, outperforming all other results on the LFW website.

Brief Method Description

The face recognition technology is comprised of Aurora’s proprietary algorithms, machine learning and computer vision techniques. We report results using the unrestricted training protocol, applied to the view 2 ten-fold cross validation test, using images provided by the LFW website, including the aligned and funnelled [6] sets and external data used solely for alignment purposes.

Company Background

Aurora has 15 years’ experience in the field of biometric face recognition. For the last few years our team has been working on the challenging problems of uncontrolled visible-spectrum colour images, as detailed in this report. We have previously created infrared face recognition systems, utilising our bespoke hardware, to overcome typical lighting problems. Our systems are widely deployed and we have more installed systems in the UK than any other competitor.

Recognition Pipeline

The complete face recognition pipeline of a full system consists of four primary steps: face detection; feature point localisation; feature extraction and classification. The verification stage (feature extraction and classification) constitutes proprietary descriptor extraction procedures and comparison metrics to produce a similarity score, which in turn is applied to a threshold for the final classification decision.

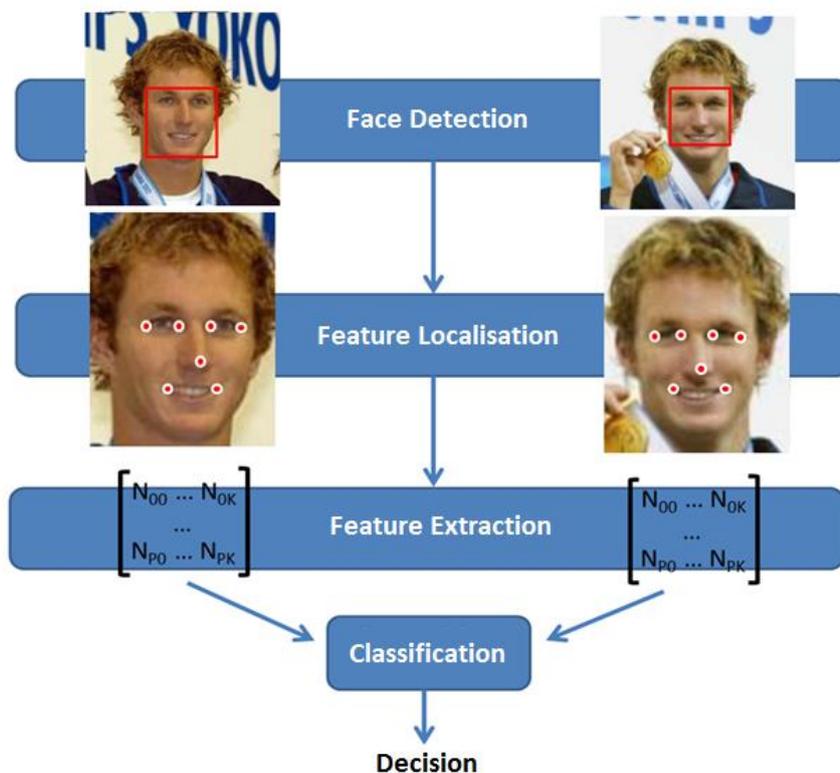


Figure 1. Aurora's recognition pipeline.

Although Aurora has developed a number of highly accurate algorithms for face detection and feature localisation (not described here), the purpose of this investigation is to evaluate the verification stage (feature extraction and classification) of the pipeline as an individual component. Therefore, we assume face detection and feature localisation have been completed successfully, relying on the pre-aligned face images.

The purpose of this investigation is to determine the maximum achievable performance, given accurately located feature points. Other experiments by ourselves investigate the impact of automatically detected feature points.

Evaluation Methodology

The evaluation is carried out according to the 10-fold cross validation test under the unrestricted configuration, strictly following the training and test procedure that was defined in the technical report of LFW database [1].

Ten 'folds' of the view 2 data set are processed. For each fold, 600 image pairs are compared to produce a similarity score, to which a threshold is applied to make the final classification decision. Images from outside of the test set pairs are used to train the face recognition model, comparison

metric and classification threshold. Although we do not restrict training to only those images pairs specified, we do ensure that no training, optimisation or fine tuning is carried out on any of the images present in the 600 image pairs of the test set. Neither are any subjects in the 600 image pairs present in the data used for training, optimisation or fine tuning; hence each fold of the test is conducted blind, as required by the protocol. The training and evaluation process is repeated for each of the ten folds, from which the mean classification accuracy is computed.

Results

Using the LFW unrestricted protocol, the Aurora face recognition engine achieves a mean classification accuracy of 0.9324 ± 0.0044 , outperforming all other results published on the LFW website.

Organisation	Algorithm	$\hat{u} \pm S_E$
Aurora	Aurora-c-2014-1	0.9324 ± 0.0044
UST China, MS Research Asia	High-dim LBP [2]	0.9318 ± 0.0107
Oxford University	Fisher vector faces [3]	0.9303 ± 0.0105
Vision Labs	VisionLabs ver.1.0, aligned	0.9290 ± 0.0031
NEC	CMD+SLBP, aligned [5]	0.9258 ± 0.0136
Face.com	Face.com r2011b [4]	0.9130 ± 0.0030

Table 1. Mean classification accuracy (\hat{u}) and standard error (S_E) of the five top performing submissions reported on the LFW website, compared with the Aurora algorithm.

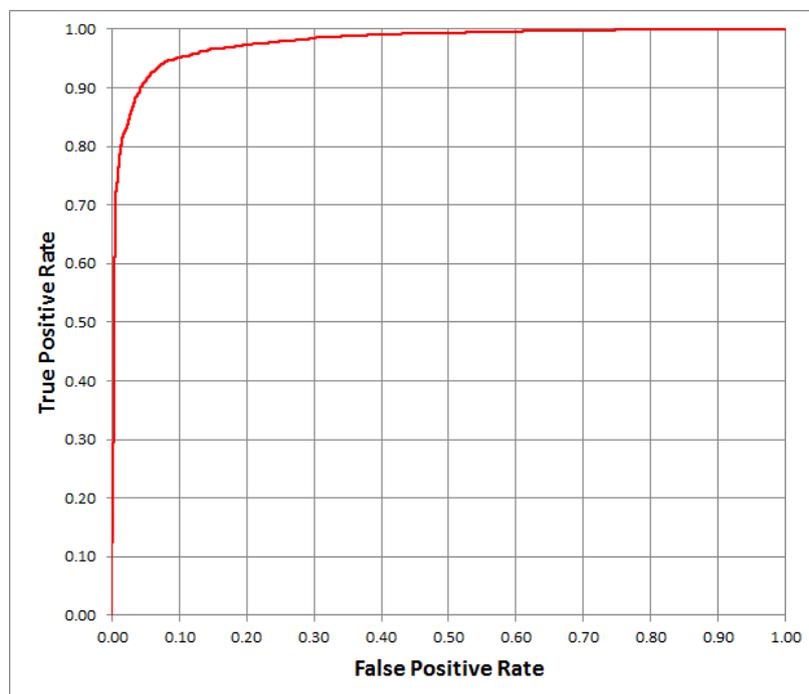


Figure 2. ROC curve of the “Aurora-c-2014-1” algorithm for the full ten-fold cross validation.

Results and Conclusion

The results produced by Aurora are state of the art and able to match faces in surprisingly difficult conditions. Example results from the LFW dataset are shown below, all of which were **correctly** classified by the Aurora system. The matches include extreme appearance changes such as different hair styles, hats and glasses, make-up and ageing, as well as more common difficulties such as pose, expression, poor lighting and partial occlusions.



Figure 3. The algorithm is able to cope with changes in head angle, even matching from a partial left profile to partial right profile.



Figure 4. The algorithm has been developed to account for natural variations in expression, as evident in the above examples.



Figure 5. The above examples demonstrate how the system is able to cope with changes in appearance, such as hair styles (top left), beards (top right), glasses (bottom left) and aging (bottom right).



Figure 6. Examples of correct classification, given some partial occlusion of the facial region.



Figure 7. Examples of extremely difficult cases, correctly classified by the Aurora face recognition algorithm under the LFW protocol, although presenting significant challenges for face and feature localisation algorithms in a complete system.

References

1. Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. University of Massachusetts, Amherst, Technical Report 07-49, October, 2007.
2. Dong Chen, Xudong Cao, Fang Wen, and Jian Sun. Blessing of Dimensionality: High-dimensional Feature and Its Efficient Compression for Face Verification. Computer Vision and Pattern Recognition (CVPR), 2013.
3. Karen Simonyan, Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Fisher Vector Faces in the Wild. British Machine Vision Conference (BMVC), 2013.
4. Yaniv Taigman and Lior Wolf. Leveraging Billions of Faces to Overcome Performance Barriers in Unconstrained Face Recognition. ArXiv e-prints, 2011.
5. Chang Huang, Shenghuo Zhu, and Kai Yu. Large Scale Strongly Supervised Ensemble Metric Learning, with Applications to Face Verification and Retrieval. NEC Technical Report TR115, 2011.
6. Gary B. Huang, Vidit Jain, and Erik Learned-Miller. Unsupervised joint alignment of complex images. International Conference on Computer Vision (ICCV), 2007.