On Combining Classifiers for Assessing Portrait Image Compliance with ICAO/ISO Standards

Markus Storer, Martin Urschler and Horst Bischof
Institute for Computer Graphics and Vision
Graz University of Technology
Inffeldgasse 16/II, 8010 Graz, Austria
{storer, urschler, bischof}@icg.tugraz.at

Josef A. Birchbauer
Siemens Biometrics Center
Siemens IT Solutions and Services
Strassgangerstrasse 315, 8054 Graz, Austria
josef-alois.birchbauer@siemens.com

Abstract: The International Civil Aviation Organization (ICAO) has selected biometrics - in particular face - as the technique for travel documents to verify the association between such documents and the person in possession of them. The International Standard Organization (ISO) has defined a standard for the digital face images to be used in Machine Readable Travel Documents. Due to the ISO/IEC 19794-5 international standard, there exists a high demand for automatically checking portrait images to assist civil service employees in decision making regarding ICAO/ISO compliance. We present a face normalization and analysis system implementing several requirements of the ISO/IEC 19794-5 specification. We address the criteria eyes-open and mouth-closed and highlight the fusion of complementary classifiers to boost performance of the overall analysis system. Our results show that classifier fusion is capable of improving the classification performance considerably as compared to a single classifier decision.

1 Introduction

Face represents one of the most commonly used biometric trait applied in many different areas, e.g., surveillance, person verification and identification. The International Civil Aviation Organization (ICAO) has selected biometrics as the technique for travel documents to verify the association between such documents and the person in possession of them [Int06]. The default biometrics used in conjunction with machine readable travel documents (MRTD) is face. To allow interoperability among systems developed by different vendors and to simplify the integration of biometric recognition in large-scale systems a standard data format for digital face images is needed. In 2004 the International Standard Organization (ISO) defined a standard for the digital face images to be used in MRTDs. The ISO/IEC 19794-5 international standard [ISO04] specifies a set of characteristics that images have to comply with as well as a record format for storing, recording and transmit-

Figure 1: Illustration of the overall image analysis system consisting of normalization (*tokenization*) and compliance analysis.

ting the facial image information.

We focus on a part of our automatic system for checking arbitrary input images for their compliance to the ISO/IEC 19794-5 specification. The overall perspective of the entire system is to provide a method that assists civil service employees in determining suitable machine readable travel document photos, thereby increasing efficiency in this selection process and significantly reducing manual work. The main challenge of an automatic system for checking ICAO/ISO compliance is its robustness to a large number of different distortions in input images regarding noise, occlusions, bad lighting situations, deviations from frontal pose, and the large variety of human faces with respect to gender, race, or appearance modifications like hair type (or lack of hair), beard, make-up and glasses. All these aspects make it particularly difficult to design a robust system for face analysis. Up to our knowledge there have been no previous publications on this topic besides [SLP+05], reporting results on an automatic face image validation system, where a number of rather simple quality aspects of face images are checked. However, a number of commercial products currently exist for ISO/IEC 19794-5 compliant face analysis. In [FFM08] a comparison of several commercial solutions is done, though the vendors associated with the presented results are kept anonymous. This work expresses the necessity in publicly available facial image benchmark data that is supplementing the formal specification of requirements.

Face and facial component detection and related analysis has a long tradition in the computer vision literature, a survey can be found in [YKA02]. Analysis of facial expressions is a hot topic in recent years [FL03], [PR00], with recognition of behavior and emotion from videos as its goal. Face tracking and assessment of medical states like fatigue or bad posture are important, e.g., in driver assistance systems [TC96] or for the prevention of work-related disorders.

Our specific interest lies in the analysis of the state of the eyes and the mouth on a given arbitrary face image. In our system first a normalization stage (which is called *tokenization*) transforms an arbitrary input image to a normalized coordinate frame depending on the eye positions by making use of a robust face and facial component detection algorithm. The overall processing pipeline is depicted in Figure 1, while the outcome of the *tokenization* process is illustrated in Figure 2.

After the detection of eyes and mouth components in a facial image, an analysis procedure is applied to assess the ICAO criteria *eyes-open* and *mouth-closed*. Reviewing the literature revealed a multitude of techniques being applied to this problem, among them machine-learning approaches like support vector machines [Vap95], or Boosting [SFBL98], model-based approaches like EigenFaces [TP91] or Active Appearance Models [CET01], or sim-

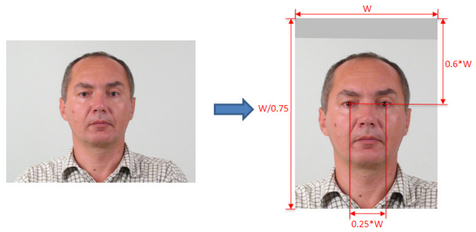

Figure 2: The tokenization procedure. The input image (left) and the tokenized image (right).

pler geometric and template based methods for detecting eye- or mouth-related features like lips, teeth, iris, or eyelids. Despite their usefulness in many situations, all of these approaches have their specific drawbacks, e.g., performance of model-based approaches decreases significantly in the presence of outliers, while their accuracy is superior to other methods if the model fitting is successful. From the observation of differing performance of different - to a certain extent complementary - algorithms, we adapted the interpretation of each algorithm as a single expert giving a vote for a certain classification decision. By combining the votes of all classifiers in a classifier fusion scheme [KHDM98] ,[Kun02], we state the hypothesis that the performance of the combined scheme is superior compared to the performance of the single classifiers in the ensemble. This makes the decision for the specific events *eyes-open* and *mouth-closed* more robust in the presence of difficult situations like noisy input data, lighting conditions or partial occlusions, e.g., wearing glasses.

The rest of the paper is organized as follows. We will validate our hypothesis by presenting our face analysis system consisting of the single classifiers (Section 2) and the classifier fusion strategies in more detail in Section 3. Section 4 shows the results of our experiments on two databases. Finally we discuss and summarize our findings in Section 5.

(a) (b) (c)

Figure 3: Some sample images from the evaluation database: (a) eyes open, mouth closed, (b) eyes open, mouth open, (c) eyes closed, mouth closed.

2 Face Analysis System Description

Our face analysis system operates on *tokenized* images. It performs several classification decisions of which we restrict ourselves to *eyes-open* and *mouth-closed* events. These criteria rely on our facial component detection stage, where a robust scheme performs face, eye and mouth localization from face component hypotheses in a probabilistic voting framework.

The ISO/IEC 19794-5 [ISO04] defines the following rules for accepting photos as suitable subject to *eyes-open* and *mouth-closed* criteria according to best practices. The face expression should be neutral (non-smiling) with both eyes open normally and the mouth closed (see Figure 3). A smile is not recommended regardless of the inside of the mouth and/or teeth being exposed or not. Starting from this specification and taking the large variety of possible problems in real-world images (due to noisy data, inappropriate lighting situations, occlusions due to hair or glasses, or the large variety in appearance of different people) it is intuitive that a single classification method will not be able to solve this task in a robust manner. Therefore several different classifiers are combined in a fusion step (see Figure 4). An important assumption for efficiently combining classifiers is that they show complementary behavior and their estimates are as independent as possible. In practice it is very hard to come up with a set of totally independent methods, so one has to rely on experimental evaluation to show their applicability to a given task.

For the training based approaches we have used a large manually annotated training set of around 4600 face images which were taken from the Caltech Face database [MS199], the FERET database [PWHR98] and from a third database constructed from our own images.

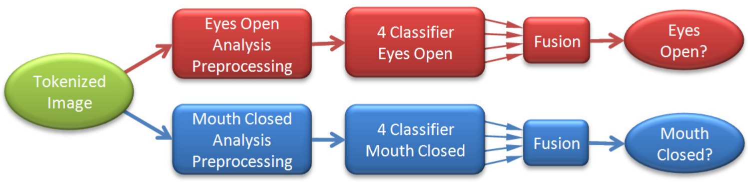

Figure 4: Face analyzer workflow. From tokenized images we perform some preprocessing, apply the single classifiers and fuse their results to form a final decision.

2.1 Eyes Open Analysis

For the analysis of the event *eyes-open* we use an ensemble of four classifiers. Two classifiers are based on AdaBoost, one uses the Active Appearance Model and the last method is based on a geometric iris localization strategy. The *eyes-open* decision is performed independently for the left and right eye and leads to a confidence value $d_i(\mathbf{x}) \in [0, 1]$ representing the range between closed and open eyes. The minimum of these two separate decisions forms the final result, since one closed eye already corresponds to an *eyes-closed* decision.

2.1.1 Active Appearance Model

We trained an Active Appearance Model (AAM) [CET01], **(CL1)**, for face image regions around the eyes, see Figure 5. Our training set consists of 427 manually annotated face images taken from the Caltech Face database and our own collection. Training images show variations in the opening of the eyes, slight pose variations, and eyes, which are looking straight and away. For model building we keep 97.5 percent of the eigenvalue energy spectrum to represent our compact model. To apply the AAM to a given image for *eyes-open* classification, we initialize the mean shape of the AAM by the roughly estimated left and right eye locations from facial component detection. To derive a measure of the likelihood of the *eyes-open* event we analyze the vertical eyes' opening of the converged AAM shape model in the left and right eye area respectively. We compare the opening to a pre-defined threshold $T_{E,aam}$ and additionally weight the distance to the threshold with the AAM residual error that represents an estimate of success or failure of model fitting.

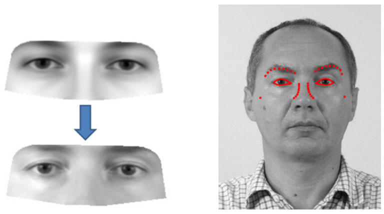

(a) (b)

Figure 5: Active Appearance Model of the eye region. (a) Learned mean shape/texture and the texture after successful fitting. (b) AAM shape model after successful fitting drawn on the input image.

2.1.2 Iris Detection Approach

Our geometric iris detection approach **(CL2)** is based on a fast radial symmetry detector presented in [LZ03]. For each eye we restrict ourselves to an image patch around the eye. After performing edge-preserving anisotropic smoothing [WtHRV98], we calculate a symmetry transform image by estimating gradient orientation and magnitude projection images over several scales according to [LZ03]. Local minima of the symmetry transform image correspond to centers of radial-symmetric structures. The strongest response of this transform corresponds to iris centers. Afterwards we perform a more accurate iris radius estimation by using a one-dimensional Hough voting on the binary response image from a Canny edge detector [Can86]. We favor iris radii that are conform with a rough scale estimation of the iris that we are able to derive from our tokenized input images. The voting histogram entry with the maximal response gives the desired iris radius. From the strength of the symmetry image minimum and the voting histogram we derive a confidence measure for the *eyes-open* event.

2.1.3 AdaBoost Classifier

Eyes-open analysis using AdaBoost [SFBL98], [VJ04], utilizes two different classifiers, both trained with the OpenCV [Int07] library. These classifiers focus on Haar wavelet filter features. The first one **(CL3)** was trained on image patches of closed eyes and the second **(CL4)** on open eyes. For the closed eye classifier 464 positive image patches were used, while the open eye classifier was trained with 2732 image patches. The discrepancy in the number of positives is due to the unequal representation of both classes in our training set. In both cases the set of negative images was taken from generic background images.

Our classification strategy for each trained Boosting classifier takes the approximate location of the eye from the facial component detection stage and applies the classifier to a slightly enlarged region around this region. That is the reason, why we trained one open-eye and one closed-eye detector utilizing the sliding window approach in the enlarged region, hence exploiting the detector as a classifier. If we detect an open eye rectangle we report a confidence measure according to [WN07]. If we detect a closed eye rectangle we report the inverse of this confidence measure.

2.2 Mouth Closed Analysis

For the *mouth-closed* analysis we have also used an ensemble of four classifiers. Three classifiers are based on AdaBoost and one classifier utilizes a blob detection algorithm that locates dark blobs due to mouth cavity shadows. Decision scores $d_i(\mathbf{x}) \in [0, 1]$ range between 0 for open and 1 for closed mouthes.

2.2.1 Geometric Dark Blob Analysis

The dark blob analysis **(CL1)** is a geometric method that makes use of the fact that open mouthes very often exhibit dark blobs due to shadows in the open mouth cavity compared to the rest of a mouth image patch. Therefore, we investigate a slightly enlarged version of the mouth detection area, transform it into HSV color space and proceed by working solely on the Value coordinate. After binarizing the mouth patch using thresholding [Ots79], we perform a blob detection process that extracts dark blobs corresponding to shadow regions. A filtering stage on the extracted blobs regarding their size, center locations and compactness removes unlikely shadow regions that may occur, e.g., due to beards. If a dark blob region survives this filtering stage we decide for the *mouth-open* event, otherwise for *mouth-closed*. A confidence measure is derived from the size of the detected blob region.

2.2.2 Boosting Classifier

Mouth-closed analysis using AdaBoost leads to three different classifiers. The first one **(CL2)** is trained with the OpenCV library using 3785 closed mouth patches as positives and a large pool of non-mouth patches as negatives. This classifier focuses on Haar wavelet filter features. The classification strategy takes the approximate location of the mouth from the facial component detection stage into account and applies the classifier to a slightly enlarged region around the mouth.

The second AdaBoost classifier **(CL4)** uses integral image approximations of edge orientation histograms. The weight update strategy follows the *RealBoost* scheme. We expect complementary behavior of our *RealBoost* approach to the OpenCV implementation due to the different features under consideration, i.e., the OpenCV library focuses on wavelet filter approximations, while our *RealBoost* learns features from the edge information of an image. *RealBoost* is applied similar to the first classifier but uses 2475 open mouth patches as positive images in the training stage.

The third AdaBoost classifier **(CL3)** is also trained with the *RealBoost* scheme on 1200 closed mouth patches. This classifier uses the same amount of open mouth patches as negatives and can only be applied directly (without a sliding window approach) to the detected mouth patches from the facial component detection. All of the AdaBoost classifiers report a confidence measure which is calculated according to [WN07] in case of closed mouths.

3 Classifier Fusion

We hypothesize that fusing multiple classifiers generates more accurate classification results compared to single classifier decisions. Hence, our goal is to evaluate different fusion strategies to combine the classifiers discussed in the previous sections. Let D denote a single classifier and $\mathbf{x} \in \mathbb{R}^n$ a feature vector representing a pattern to be classified. The classifier represents a mapping

$$D:\ \mathbf{x} \in \mathbb{R}^n \rightarrow \omega_j \in \Omega\,,$$

where ω_j is one of the c possible classes of $\Omega = \{\omega_1, \ldots, \omega_c\}$. Denote $\{D_1, \ldots, D_L\}$ as the set of L classifiers. The output of the ith classifier is $D_i(\mathbf{x}_i) = [d_{i,1}(\mathbf{x}_i), \ldots, d_{i,c}(\mathbf{x}_i)]^T$, where $\mathbf{x}_i$ is the specific feature vector representation of the input pattern needed by classifier D_i and $d_{i,j}(\mathbf{x}_i)$ is the confidence, i.e., the degree of support, classifier D_i assigns to the assumption of $\mathbf{x}_i$ originating from class j. The fused output $\hat{D}$ of the L single classifiers is

$$\hat{D}(\mathbf{x}) = \mathcal{F}(D_1(\mathbf{x}), \ldots, D_L(\mathbf{x})), \tag{1}$$

where $\mathcal{F}$ is called the *fusion strategy*. Resulting from $\hat{D}$, the final confidence values assigned to each class are $\hat{d}_j$.

The following fusion strategies are investigated:

Minimum (MIN) $\hat{d}_j(\mathbf{x}) = \min\limits_i\{d_{i,j}(\mathbf{x})\}$

Maximum (MAX) $\hat{d}_j(\mathbf{x}) = \max\limits_i\{d_{i,j}(\mathbf{x})\}$

Average (AVR) $\hat{d}_j(\mathbf{x}) = \frac{1}{L}\sum\limits_{i=1}^{L} d_{i,j}(\mathbf{x})$

Binarized Average (BAVR) This scheme is equivalent to the average fusion strategy, except for the outputs of the single classifiers $d_{i,j}$ being assigned to a specific class explicitly before averaging.

Product (PRO) $\hat{d}_j(\mathbf{x}) = \prod_{i=1}^{L} d_{i,j}(\mathbf{x})$

Prior Confidence (PRIOR) A priori confidences of the single classifiers are obtained from tests on a validation dataset according to their performance exhibited (Table 1). The prior confidences are accumulated according to the decision of the corresponding single classifier.

Bayes Combination (BAYES) This scheme assumes that the classifiers are mutually independent and that the posterior confidences of the single classifiers are equal to posterior probabilities. For our single classifiers we expect independence, because different underlying concepts and methodologies are used, e.g., different features for classification.

The fusion strategies presented above can only be justified under strict probabilistic conditions. Nevertheless, some of them exhibit excellent performance as can be seen in the experimental results, a fact which was already stated in [KHDM98] for some of the classifier fusion strategies.

4 Experimental Results

We used two different face databases for the evaluation of our single classifiers and fusion strategies. The first one is the publicly available **AR face database** [MB98]. It consists of 126 unique people, frontal view face images (70 male, 56 female) of different facial expressions, illumination conditions and occlusions resulting in a total amount of over 4000 color images of a resolution of 768 by 576. We used all the images except those with dark sunglasses or occluded mouth due to a scarf yielding about 1700 images for evaluation. Annotation data is available on request. Some samples of this challenging database are given in Figure 6. The **second database** we used is a private data set containing 325 frontal face color images of 480 by 640 pixels with 30 people showing different facial expressions.

All evaluation results on the AR face database are summarized in Table 2 for *eyes-open* and for *mouth-closed* analysis. The evaluation results on our private database are presented in Table 3. The comparison of the ROC curves of the best single classifier to the best fusion strategy is shown in Figure 7 for the AR face database and in Figure 8 for our own database.

The best overall fusion performance is exhibited by the Bayes combination (BAYES) strategy. The single classifier performance on our private database is slightly better compared to the AR database due to a larger complexity of the images contained in the latter data set. The AR face database contains a large number of images from people wearing different glasses, often show severe specular reflections, simulate bad illumination conditions and

Figure 6: Sample images from the AR face database [MB98] showing the difficulties of the images under consideration.

Table 1: Prior confidences of the single classifiers. (left) eyes-open classifiers, (right) mouth-closed classifiers

Classifier	Prior
CL1	0.25
CL2	0.25
CL3	0.35
CL4	0.15

Classifier	Prior
CL1	0.23
CL2	0.21
CL3	0.25
CL4	0.31

extremely wide open mouths, where we have problems of robustly locating the mouth region. However, the hypothesis of fusing multiple classifiers generating more accurate and robust classification results compared to a single classifier, is thus approved, illustrated in our figures and tables.

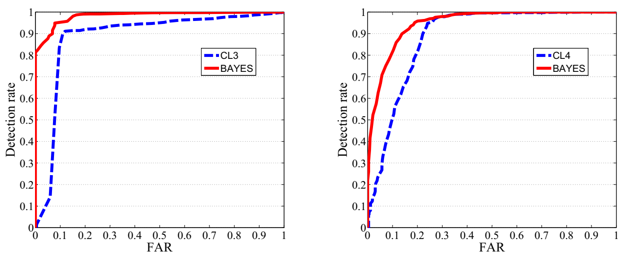

(a) (b)

Figure 7: Comparison of the ROC curves of the best single classifier to the best fusion strategy for the (a) eyes-open and (b) mouth-closed analysis on the AR face database.

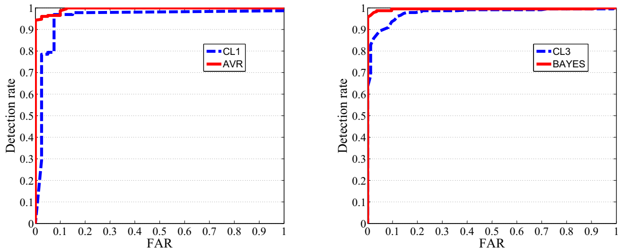

(a) (b)

Figure 8: Comparison of the ROC curves of the best single classifier to the best fusion strategy for the (a) eyes-open and (b) mouth-closed analysis on our own face database.

Table 2: Evaluation results in [%] for the eyes open analysis (left) and for the mouth-closed analysis (right) on the AR face database

	EER	Detection Rate @ FAR			
		5	10	30	50
CL1	18.35	48.90	68.68	86.66	94.14
CL2	29.37	70.63	70.63	70.63	70.63
CL3	10.85	11.96	85.18	93.49	95.04
CL4	24.89	74.45	74.61	75.37	76.65
MIN	25.35	73.09	73.49	75.00	76.48
MAX	16.00	29.00	61.47	98.89	99.70
AVR	8.48	90.31	91.82	97.03	99.80
BAVR	9.45	90.55	90.55	90.55	99.59
PRO	20.65	78.55	78.80	79.83	80.85
PRIOR	10.40	89.60	89.60	99.05	99.59
BAYES	7.35	87.65	95.26	99.22	99.69

	EER	Detection Rate @ FAR			
		5	10	30	50
CL1	36.90	0.00	0.00	0.00	92.31
CL2	28.33	64.83	66.03	72.18	80.17
CL3	25.13	31.49	44.88	79.67	95.31
CL4	19.53	25.34	50.80	97.83	99.67
MIN	18.91	61.38	66.76	90.56	92.04
MAX	37.86	8.06	8.12	8.35	99.75
AVR	12.99	67.84	80.85	97.52	99.67
BAVR	14.91	50.63	50.63	90.23	99.00
PRO	14.29	66.92	78.54	95.71	97.25
PRIOR	13.96	50.63	50.63	96.24	99.25
BAYES	12.56	66.23	81.55	97.92	99.67

Table 3: Evaluation results in [%] for the eyes open analysis (left) and for the mouth-closed analysis (right) on our own face database

	EER	Detection Rate @ FAR			
		5	10	30	50
CL1	7.50	79.39	96.93	97.92	98.14
CL2	10.96	89.04	89.04	89.04	89.04
CL3	13.51	85.09	85.53	92.65	94.88
CL4	12.72	87.28	87.28	87.28	87.28
MIN	13.16	86.84	86.84	96.38	97.04
MAX	6.60	86.40	99.56	100.00	100.00
AVR	3.95	96.49	98.68	100.00	100.00
BAVR	3.95	96.05	96.05	100.00	100.00
PRO	8.97	90.92	91.05	91.58	92.11
PRIOR	5.00	99.56	99.56	100.00	100.00
BAYES	5.00	97.37	98.03	99.56	99.78

	EER	Detection Rate @ FAR			
		5	10	30	50
CL1	12.02	87.98	87.98	87.98	87.98
CL2	11.86	53.48	85.49	94.97	95.78
CL3	8.62	89.36	93.82	98.71	99.14
CL4	12.39	55.56	81.93	98.02	99.57
MIN	12.02	87.98	87.98	87.98	87.98
MAX	6.02	22.75	99.27	99.57	99.57
AVR	2.15	98.35	99.14	99.57	99.57
BAVR	4.29	95.71	98.28	98.28	99.57
PRO	3.26	97.85	97.85	98.08	98.51
PRIOR	4.29	97.00	98.28	99.57	99.57
BAYES	2.28	98.71	99.57	99.57	99.57

5 Conclusion

In this work we present selected parts of our portrait analysis system for checking ISO/IEC 19794-5 compliance. ICAO/ISO requirements on portrait images for machine readable travel documents need to be assessed automatically for issuing documents like e-passports, identification cards, or visa. Within this paper we specifically deal with the two criteria eyes-open and mouth-closed that require high-level facial component understanding. Both criteria are challenging due to noise, undesired lighting, occlusions (glasses, beard) and the large variety in human faces in itself. We overcome these difficulties by a fusion of complementary classification methods. Within a classifier fusion framework we are testing a number of different fusion strategies to combine the votes of single classifiers. Our experimental results show that classifier fusion is capable of improving the classification performance considerably, thus validating our hypothesis. The best overall fusion performance is exhibited by the Bayes combination (BAYES). Based on our findings, further work is necessary to evaluate additional, complementary classification schemes to further improve the overall classification results for both criteria, as well as transferring our approach to portrait image criteria additionally covered in ISO/IEC 19794-5 (e.g., "deviation from frontal pose", "eyes looking away").

Acknowledgement

This work has been funded by the Biometrics Center of Siemens IT Solutions and Services, Siemens Austria.

References

[Can86] J. Canny. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[CET01] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active Appearance Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):681–685, 2001.

[FFM08] M. Ferrara, A. Franco, and D. Maltoni. Evaluating Systems Assessing Face-Image Compliance with ICAO/ISO Standards. In *Proceedings 1st European Workshop on Biometrics and Identity Management (BIOID)*, May 2008.

[FL03] Beat Fasel and Juergen Luettin. Automatic Facial Expression Analysis: A Survey. *Pattern Recognition*, 36(1):259–275, 2003.

[Int06] International Civil Aviation Organization (ICAO). Doc 9303, Machine Readable Travel Documents, Part 1 - Machine Readable Passport, Sixth Edition, 2006.

[Int07] Intel. Open Computer Vision (OpenCV) Library. http://www.intel.com/technology/computing/opencv/index.htm, 2007.

[ISO04] ISO International Standard ISO/IEC JTC 1/SC37 N506. Biometric Data Interchange Formats - Part 5: Face Image Data, 2004.

[KHDM98] Josef Kittler, Mohamad Hatef, Robert P.W. Duin, and Jiri Matas. On Combining Classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):226–239, March 1998.

[Kun02] L. Kuncheva. A Theoretical Study on Six Classifier Fusion Strategies. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(2):281–286, 2002.

[LZ03] Gareth Loy and Alexander Zelinsky. Fast Radial Symmetry for Detecting Points of Interest. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8):959–973, August 2003.

[MB98] A.M. Martinez and R. Benavente. The AR Face Database. Technical Report 24, CVC, June 1998.

[MS199] Caltech Face Database. http://www.vision.caltech.edu/html-files/archive.html, 1999.

[Ots79] N. Otsu. A threshold selection method from gray-level histogram. *IEEE Transactions on System Man Cybernetics*, SMC-9(1):62–66, 1979.

[PR00] Maja Pantic and Leon J. M. Rothkrantz. Automatic Analysis of Facial Expressions: The State of the Art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12):1424–1445, December 2000.

[PWHR98] P.J. Phillips, H. Wechsler, J. Huang, and P. Rauss. The FERET Database and Evaluation Procedure for Face Recognition Algorithms. *Image and Vision Computing*, 16(5):295–306, 1998.

[SFBL98] R. E. Schapire, Y. Freund, P. Bartlett, and W. S. Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *Ann. Stat.*, 26(5):16511686, 1998.

[SLP$^+$05] M. Subasic, S. Loncaric, T. Petkovic, H. Bogunovic, and V. Krivec. Face Image Validation System. In *Proceedings 4th International Symposium on Image and Signal Processing and Analysis*, pages 30–33, 2005.

[TC96] David Tock and Ian Craw. Tracking and Measuring Drivers' Eyes. In *Proceedings 7th British Machine Vision Conference*, volume 14, pages 541–547, August 1996.

[TP91] Matthew Turk and Alex Pentland. Eigenfaces for Recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.

[Vap95] V. N. Vapnik. *The nature of statistical learning theory*. Springer-Verlag New York, 1995.

[VJ04] Paul Viola and Michael J. Jones. Robust Real-Time Face Detection. *International Journal of Computer Vision*, 57(2):137–154, 2004.

[WN07] Bo Wu and Ram Nevatia. Improving Part based Object Detection by Unsupervised, Online Boosting. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, pages 918–927, June 2007.

[WtHRV98] J. Weickert, B. M. ter Haar Romeny, and M. A. Viergever. Efficient and Reliable Schemes for Nonlinear Diffusion Filtering. *IEEE Transactions on Image Processing*, 7(3):398–410, 1998.

[YKA02] Ming Hsuan Yang, David Kriegman, and Narendra Ahuja. Detecting Faces in Images: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(1):34–58, January 2002.