Local Facial Asymmetry for Expression Classification[*]

Sinjini Mitra
smitra@stat.cmu.edu
Department of Statistics
Carnegie Mellon University

Yanxi Liu
yanxi@cs.cmu.edu
The Robotics Institute
Carnegie Mellon University

Abstract

We explore a novel application of facial asymmetry: expression classification. Using 2D facial expression images, we show the effectiveness of automatically selected local facial asymmetry for expression recognition. Quantitative evaluations of expression classification using local asymmetry demonstrate statistically significant improvements over expression classification results on the same data set without explicit representation of facial asymmetry. A comparison of discriminative local facial asymmetry features for expression classification versus human identification is given.

1 Introduction

Human faces have an approximate bilateral symmetry. This fact has been exploited by both psychologists and computer scientists. On the one hand, human faces are assumed to be *symmetrical*. Under this assumption, Zhao and Chellappa [17], and Shimshoni *et al* [15] developed feasible 3D face reconstruction algorithms; and Seitz and Dyer achieved visually believable 3D morphing of Mona Lisa's face [14]. On the other hand, human facial asymmetry has long been acknowledged and is considered as a critical factor in attractiveness [16] and expressions [3, 4, 7, 13]. For human identification purposes, half faces have been tried, with some reported differences in recognition rates [12] while others did not [8]. Liu *et al* have shown that quantified facial asymmetry contains discriminating information for human identification under expression variations and statistically significant improvements can be made by combining facial asymmetry with classical human identification algorithms [11]. More recently, Dovgard and Basri [6] observed and recorded a correlation between worse 3D face reconstruction results with increased facial asymmetry of the subjects in the image. Liu and Palmer have demonstrated that the existence of facial asymmetry on a 3D image data set of 111 subjects is statistically significant (p < 0.001) and developed a gender classi-

fier (96% ± 2.3% classification rate) using quantified, local facial asymmetry alone [10].

Expression recognition has a direct relation to face recognition [12]. In both [11] and [12], the authors show quantitatively convincing results that face recognition rates depend on different types of facial expressions; which echoes human behavior as well. In addition, facial expression recognition plays an important role in human computer interaction.

The results of [11] suggest that facial asymmetry measures are biometrics that are potentially "invariant" to human expression variations. In this paper, we investigate (1) whether facial asymmetry as defined in [11] contains discriminative information for expression classification; (2) whether local facial asymmetry enhances the performance of an automatic expression classification algorithm; and (3) where the automatically selected local facial asymmetry features reside on human faces for expression versus human identity classifications. Though expression classification systems have been reported (e.g. [5]), to the best of our knowledge, facial asymmetry has not been used for automatic expression classification and its discriminativeness has not been analyzed against that in human identification.

Figure 1: Normalized faces of six subjects from Cohn-Kanade AU-Coded Facial Expression Database [9] with joy (top row), disgust (middle row) and anger (bottom row) expressions.

2 Experimental Data

We use the same facial expression data, same normalization process and the same definition for facial asymmetry as reported in [11]. For clarity, we repeat the data description in this section.

The dataset consists of frontal views of human expressions under controlled bilaterally balanced lighting (Figure 1). 55 subjects are used in our experiments for expression classification. Each subject has 3 expressions: joy, anger and disgust. Each face image is normalized by an affine transformation using three points: left and right inner canthi and the philtrum¿ Two types of facial asymmetries of a normalized face image are defined:

Density Difference _D-face_:

$$D(x, y) = I(x, y) - I'(x, y) \qquad (1)$$

Edge Orientation Similarity _S-face_:

$$S(x, y) = \cos(\phi_{I_e(x,y), I'_e(x,y)}) \qquad (2)$$

where $\phi_{I_e(x,y), I'_e(x,y)}$ is the angle between the edge orientations of two edge images I_e, I'_e at the same pixel location (x, y). See Figure 2 for a sampler of computed asymmetry faces under this definition. We
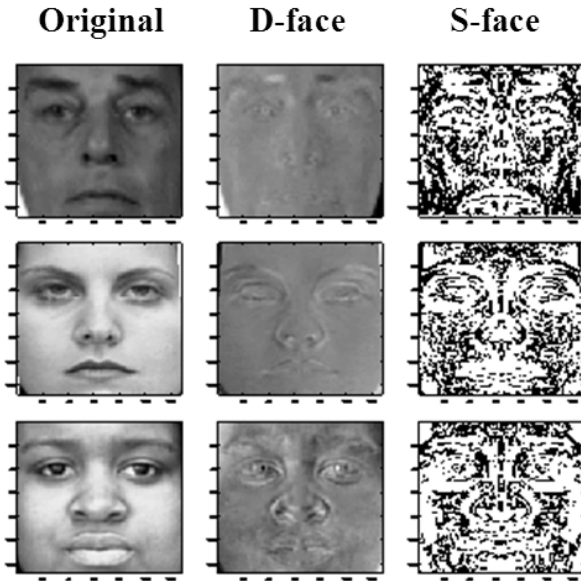

Figure 2: Some samples of original face, Difference Face (D-Face) and Similarity Face (S-Face) under the same definition as used in [11].

denote these half faces as _D-face_ and _S-face_ which, together with their various projections (Table 1), are called _AsymmetryFaces_.

Table 1: A set of AsymmetryFaces are defined. The data sets used for face recognition experiments contain 55 subjects, each has three expressions. D-face and S-face are represented by the top 60 and 100 principal components respectively, capturing the 95% variance in the data.

Data	Description	Features
D	top 60 Eigen vectors of _D-face_	60
D_x	column-mean of _D-face_ on X axis	64
D_y	row-mean of _D-face_ on Y axis	128
S	top 100 eigen vectors of _S-face_	100
S_x	column-mean of _S-face_ on X axis	64
S_y	row-mean of _S-face_ on Y axis	128

3 Expression Identification

The motivation for using facial asymmetry for expression classification comes from observations in the literature that facial expressions have a strong relation with facial asymmetry, and from the reported results that quantified facial asymmetry measures are useful for human identification under expression variations. Our goal in this study is to find out whether AsymmetryFaces defined above (Table 1) contain expression discriminating information.

This is a classification problem with 3 classes: joy, anger and disgust (Figure 1). The experimental setup starts with a random division of 55 subjects into 30 and 25 sized subsets using the size-30-subset for training and the size-25-subset for testing. This division process is repeated twenty times. The final results (False Positive Rates, False Negative Rates) are obtained by averaging over these 20 repetitions. We have chosen a FisherFace [1] classifier (even though our classes are expressions instead of human identities, as the classifier was originally intended), as a baseline classifier for comparison. A sequential forward feature selection (SFS) algorithm combined with an Linear Discriminant Analysis (LDA) Algorithm is used for evaluation of final results. Therefore, the results contain three parts: (1) the results from FisherFace classifier; (2) the results from AsymmetryFaces only using a SFS + LDA classifier; and (3) the results from applying SFS + LDA on features from both FisherFace and AsymmetryFaces.

Table 2 shows the results of expression identification using FisherFace on original images, using AsymmetryFaces alone, and FisherFace + AsymmetryFaces. Some observations:

- The classification rate of Fisherface classifier is

about 85% using normalized facial information with an average false negative rate (FNR) of 14.6% and false positive rate (FPR) of 7.3%.

- While using each AsymmetryFace alone, the FNRs are much higher than those from the FisherFace: 17.8% to 47.67%, the combined AsymmetryFaces, especially D-faces with S-faces, improve the classification rates. The combination of all AsymmetryFaces provides a comparable expression classification rates (average FNR 13.2% and FPR 6.6%) and smaller variance than those from the FisherFace.

- Combining AsymmetryFaces with eigen features used in the FisherFace classifier yields an effective improvement on the classification rates, the FNR drops from 14.6% to 3.6%, and FPR from 7.3% to 1.8%.

- In general, S-face and its projections show more discriminative power than D-face and its projections, with lower FNRs. The role of D-face and S-face for human identification is just the opposite as reported in [11]. Here S-face and its projections have a higher average classification rate.

- Finally, the joy expression seems to be easier to distinguish with both a lower FNR and a lower FPR than the other two expressions. It is worth noting that in [11], face recognition with joy expression is the most difficult (highest error rate).

3.1 Significance Evaluation

From Table 2, one can observe that the results of the FisherFace alone and combined AsymmetryFaces alone are not significantly different from each other at both the 1% and 5% confidence levels (p-value of 0.2272). Table 2 shows a reduction in error rates when facial asymmetry information is added. Furthermore, there is significant statistical difference between the results from the FisherFace alone and those from the combination of FisherFace and AsymmetryFaces (p-value of nearly 0 from the paired t-test).

3.2 Local Facial Asymmetry Analysis

The effectiveness of facial asymmetry for expression classification reported in this paper may appear to be non-intuitive given the results on the same data set [11] for human identification under expression variation. Generally speaking, when a feature (e.g. facial asymmetry) is discriminative under expression variations, it is logical to assume that this feature remains relatively invariant to different expressions and thus is not a good candidate discriminative feature for expression classification. The key observation in this work

is that it is the *location* of facial asymmetry features that contributed differently to different classification problems (human identification versus expression classification). These locations are automatically selected using feature selection algorithms. Feature selection is a common technique [2] for space reduction. Here we use it as a way to locate discriminative facial asymmetry features and understand how different local facial asymmetry features contribute to both human identification under expression variation and expression classification across different subjects respectively.

In our experiments, we evaluated each dimension of an asymmetry face for its discriminative value for differentiating expressions. There are many measures that can be used to evaluate the discriminating power of each individual feature dimension. To be consistent with LDA, we choose variance ratios. It is possible for a feature to have small within-class scatter for a class even if the mean value of the feature for that class is very close to the mean value of the feature in another class. Thus we define an augmented variance ratio (AVR) as follows

$$AVR(F) = \frac{Var(S_F)}{\frac{1}{C}\sum_{i=1..C}\frac{Var_i(S_F)}{min_{i\neq j}(|mean_i(S_F)-mean_j(S_F)|)}}$$

where $mean_i(F)$ is the mean of feature F's values in class i. AVR is the ratio of the variance of the feature between subjects to the variance of the feature within subjects, with an added penalty for features which may have small intra-class variance but which have close inter-subject mean values. The features that have higher variance ratios are more discriminative. Figure 3 shows the AVR value of each asymmetry feature in terms of its discriminative power for expression classification.

3.3 Comparing Local Asymmetry Features for Expression and Human Classification

Using the same AVR evaluation function also provides a convenient way to compare with the facial asymmetry dimensions selected for human identification tasks carried out in [11].

Figure 4 shows the overlaid AVR plots for each of the two projections of D-face (D_x, D_y) and S-face (S_x, S_y) and for both human identification (borrowed from [11]) and expression classification problems. These plots describe precisely which local asymmetry features from which specific part of the face are useful for which classification problems. Obviously, the features with high discriminative power for human identification have low AVR values for expression classification, and vice versa. Different dimensions of
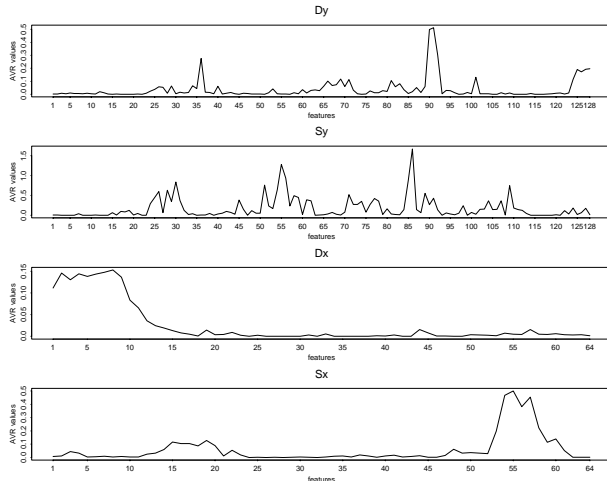
Figure 3: The AVR values for each feature dimension of the AsymmetryFaces (Table 1) for the classification problem of expression identification across different subjects.

the same type of facial asymmetry features are chosen for different purposes. The anatomical locations of the most discriminative features for identity versus expression classification goals are nearly mutually exclusive (Figure 4).

4   Discussion and Conclusion

In this novel application of facial asymmetry to expression classification, we have purposefully used a very simple and well-understood feature selection scheme and classifier to validate the utility of facial asymmetry features. Our experimental results show that quantified facial asymmetry features can perform well in expression identification across subjects in addition to their reported synergy in human identification under expression variations [11]. Our experimental results show that facial asymmetry has sufficient discriminating power to improve the performance of an expression classifier significantly.

Figure 4 provides some insight into why this is the case. Basically, the two overlaid AVR plots of facial asymmetry features for the two different classification problems guide the feature selection process to choose different local dimensions to discriminate human identity or expression. One can easily observe that features that are considered most discriminative in each case differ from each other. Thus for each classification, a different feature subspace is used. Therefore, there is no contradiction in using the same type of features, but at different locations, to achieve good results for seemingly conflicting classification goals.

References

[1] P.N. Belhumeur, J.P. Jespanha, and D.J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *PAMI*, 19(7):711–720, July 1997.

[2] C. M. Bishop. *Neural Networks for Pattern Recognition.* Clarendon Press, 1995. ISBN:0198538499.

[3] J.D. Borod, E. Koff, S. Yecker, C. Santschi, and J.M. Schmidt. Facial asymmetry during emotional expression: Gender, valence and measurement technique. *Psychophysiology*, 36(11):1209–1215, 1998.

[4] R. Campbell. The lateralization of emotion: A critical review. *International Journal of Psychology*, 17:211,219, 1982.

[5] M. Cohen, J. Shade, S. Hiller, and O. Deussen. Wang tiles for image and texture generation. In *SIGGRAPH*, pages 287–294, 2003.

[6] R. Dovgard and R. Basri. Statistical symmetric shape from shading for 3d structure recovery of faces. In *European Conference on Computer Vision (ECCV'04)*, May 2004.

[7] L.G. Fakas. Facial asymmetry in healthy north american caucasians. *Angle Orthodontist*, 51(1):70–77, 1981.

[8] S. Gutta, V. Philomin, and M. Trajkovic. An investigation into the use of partial-faces for face recognition. In *International Conference on Automatic Face and Gesture Recognition*, pages 33,38, Washington, D.C., May 2002. IEEE Computer Society.

[9] T. Kanade, J.F. Cohn, and Y.L. Tian. Comprehensive database for facial expression analysis. In *4th IEEE International Conference on Automatic Face and Gesture Recognition*, Grenoble, March 1999. Publically available at http://www.ri.cmu.edu/projects/project_420.html.

[10] Y. Liu and J. Palmer. A quantified study of facial asymmetry in 3d faces. In *IEEE International Workshop on Analysis and Modeling of Faces and Gestures*, pages 222–229. IEEE, 2003.

[11] Y. Liu, K. Schmidt, J. Cohn, and S. Mitra. Facial asymmetry quantification for expression invariant human identification. *Computer Vision and Image Understanding Journal*, 91(1/2):138–159, Special issue on face recognition, Martinez, Yang and Kriegman (Eds.). July/August 2003.

[12] A.M. Martinez. Recognizing imprecisely localized, partially occluded and expression variant faces from a single sample per class. *IEEE Transactions on Pattern analysis and machine intelligence*, 24(6):748–763, 2002.

[13] C.K. Richardson, D. Bowers, R.M. Bauer, K.M. Heilman, and C.M. Leonard. Digitizing the moving face during dynamic displays of emotion. *Neuropsychologia*, 38(7):1028–1039, 2000.

[14] S.M. Seitz and C.R. Dyer. View morphing. *SIGGRAPH*, pages 21–30, 1996.

[15] I. Shimshoni, Y. Moses, and M. Lindenbaum. Shape reconstruction of 3d bilaterally symmetric surfaces. *International Journal of Computer Vision*, 39(2):97–100, 2000.

[16] R. Thornhill and S. W. Gangestad. Facial attractiveness. *Trans. in Cognitive Sciences*, 3(12):452–460, December 1999.

[17] W.Y. Zhao and R. Chellappa. Symmetric shape-from-shading using self-ratio image. *IJCV*, 45(1):55–75, October 2001.

Table 2: Test results for **Expression Classification** in terms of average false negative rate (FNR), false positive rate (FPR) and standard deviations (STD) over 20 repetitions.

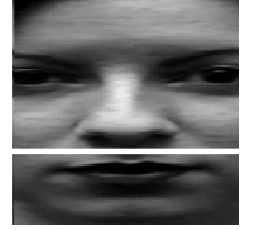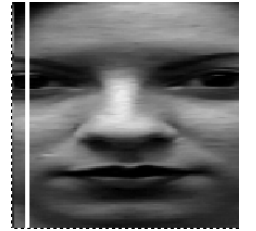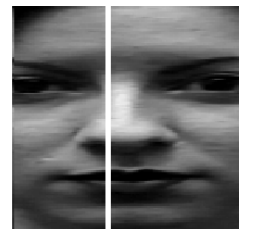| Dataset	Error Rates	Joy	Anger	Disgust	Average	STD
FF FisherFace alone	FNR	6.80%	19.40%	17.60%	14.60%	9.89%
	(STD)	(5.21%)	(9.82%)	(7.72%)	(3.63%)	(4.21%)
	FPR	2.30%	8.30%	11.30%	7.30%	6.27%
	(STD)	(2.70%)	(5.16%)	(5.44%)	(1.82%)	(2.75%)
D AsymmetryFace alone	FNR	25.80%	42.80%	41.60%	36.73%	13.32%
		(9.67%)	(11.17%)	(10.25%)	(3.88%)	(6.85%)
	FPR	19.80%	17.20%	18.10%	18.36%	6.57%
		(7.42%)	(6.88%)	(5.05%)	(1.94%)	(3.83%)
S	FNR	15.60%	17.80%	20.00%	17.80%	7.93%
		(8.50%)	(8.15%)	(9.97%)	(3.88%)	(5.95%)
	FPR	7.10%	8.10%	11.50%	8.90%	5.07%
		(5.86%)	(3.46%)	(4.94%)	(1.94%)	(2.88%)
D_Y	FNR	44.60%	37.00%	37.20%	39.60%	10.83%
		(8.03%)	(13.29%)	(9.28%)	(2.74%)	(6.98%)
	FPR	16.30%	21.00%	23.00%	19.8%	7.41%
		(6.23%)	(7.12%)	(6.54%)	(1.37%)	(4.17%)
S_Y	FNR	14.00%	20.80%	21.20%	18.67%	7.25%
		(5.58%)	(7.06%)	(7.00%)	(3.22%)	(3.21%)
	FPR	8.30%	8.10%	11.60%	9.33%	4.14%
		(4.22%)	(3.34%)	(4.33%)	(1.66%)	(2.40%)
D_X	FNR	37.20%	57.60%	48.20%	47.67%	13.01%
		(9.37%)	(9.39%)	(10.50%)	(3.59%)	(7.44%)
	FPR	26.00%	18.10%	27.10%	23.83%	9.10%
		(7.40%)	(6.57%)	(8.30%)	(1.79%)	(3.88%)
S_X	FNR	39.20%	51.80%	41.40%	44.13%	11.62%
		(10.27%)	(10.26%)	(10.96%)	(3.74%)	(7.07%)
	FPR	24.00%	19.00%	23.20%	22.07%	9.15%
		(10.76%)	(6.94%)	(9.63%)	(1.87%)	(6.53%)
S_Y+D_Y	FNR	12.20%	19.00%	18.60%	16.60%	5.92%
		(5.10%)	(5.93%)	(5.55%)	(2.72%)	(3.65%)
	FPR	8.00%	7.30%	9.60%	8.30%	3.89%
		(3.18%)	(3.80%)	(3.65%)	(1.36%)	(1.29%)
S_Y+S_X	FNR	13.80%	21.40%	17.40%	17.53%	7.38%
		(6.15%)	(6.39%)	(7.92%)	(3.57%)	(3.08%)
	FPR	8.10%	7.00%	11.20%	8.77%	4.22%
		(3.81%)	(4.13%)	(4.27%)	(1.79%)	(2.52%)
$S+S_Y+S_X$	FNR	14.10%	16.10%	13.80%	15.07%	6.32%
		(8.20%)	(6.56%)	(6.40%)	(3.62%)	(3.79%)
	FPR	5.40%	6.85%	9.75%	7.53%	4.48%
		(4.64%)	(3.10%)	(4.94%)	(1.81%)	(2.62%)
$S+S_Y+D+D_Y+S_X+D_X$ (Combine all AsymmetryFaces)	FNR	11.60%	14.60%	13.40%	13.20%	6.07%
		(7.33%)	(5.07%)	(7.14%)	(3.43%)	(3.35%)
	FPR	4.90%	6.20%	8.70%	6.60%	3.85%
		(3.70%)	(2.50%)	(4.41%)	(1.71%)	(2.53%)
$FF+S+S_Y+S_X+D+D_Y+D_X$ (Combine FisherFace with AsymmetryFaces)	FNR	1.80%	3.80%	5.20%	3.60%	3.82%
	(STD)	(3.55%)	(3.03%)	(5.21%)	(1.94%)	(2.55%)
	FPR	0.30%	2.60%	2.50%	1.80%	2.26%
	(STD)	(0.98%)	(2.44%)	(2.42%)	(0.97%)	(1.16%)
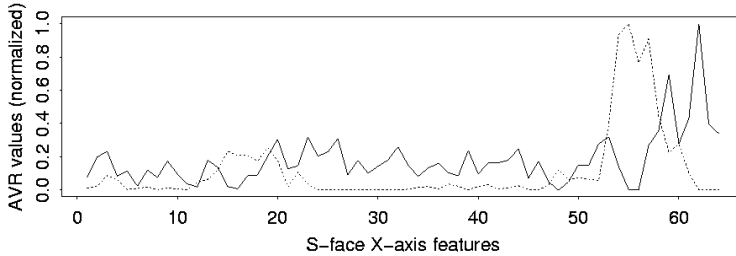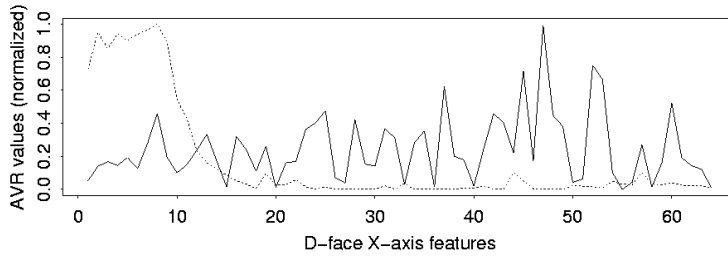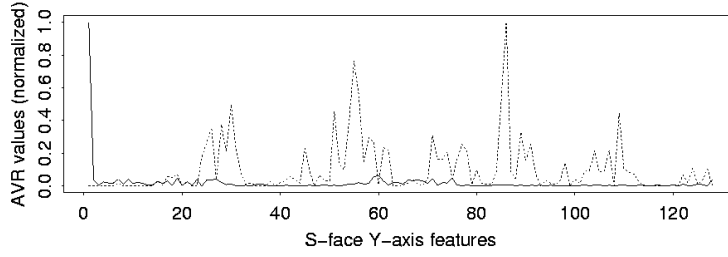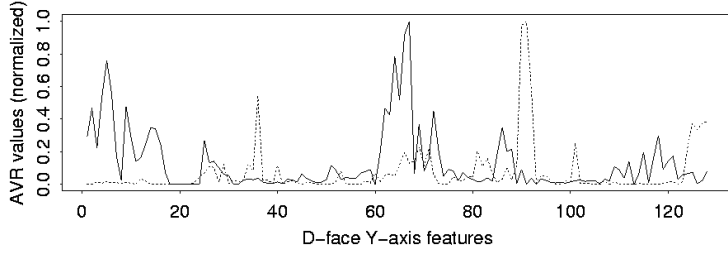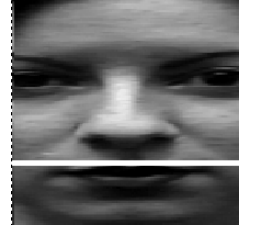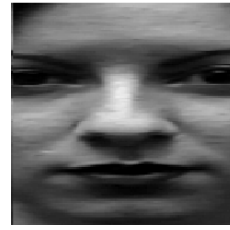
Figure 4: *Each plot (left column) shows the AVR value of each of the AsymmetryFaces (D_x, D_y, S_x, S_y defined in Table 1). The higher the AVR value for a single dimension the more discriminating power that individual feature has. D_y and S_y (128 dimensions) start from forehead (left) to chin (right). D_x and S_x (64 dimensions) start from side face (left) towards the face midline (right). (A) Human identification under expression variations. (B) Expression identification across different subjects. One can observe that the features with high AVR values for problem (A) have low AVR values for classification problem (B), and vice versa. This implies that different subsets of facial asymmetry features are selected for the two different discrimination problems. The white lines on face images of the right two columns indicate the anatomical location of the facial asymmetry feature with the highest AVR value. Note: S_y's highest AVR value for human identification is on the top of the face.*