

Is There a Connection Between Face Symmetry and Face Recognition?

Josh Harguess and J. K. Aggarwal
Computer & Vision Research Center / Department of ECE
The University of Texas at Austin
{harguess, aggarwaljk}@mail.utexas.edu

Abstract

Recent research in the area of automatic machine recognition of human faces has shown that there may be an advantage in utilizing face symmetry to improve recognition accuracy. While promising, this work has led to several open questions. What is a good feature description or score of the symmetry of the face? Is there a statistical significance between face symmetry and face recognition? We present new symmetry scores of the face and use the scores to compare the symmetry in several subgroups of a face database. A 3D face database is used to remove the effects of illumination which should improve the reliability of the symmetry score. We find a significant difference in face symmetry between the men and women subjects in the database. The database is then partitioned into most symmetric and least symmetric subjects based on the symmetry scores. The average-half-face is utilized in our face recognition experiments to take into account the symmetry of the face. Face recognition with eigenfaces using the average-half-face is significantly higher than using the full face in all subgroups regardless of symmetry score. However, face recognition using the full face does depend on the symmetry score and generally favors the least symmetric subjects.

1. Introduction

Face recognition is an important area of computer vision research and has gained significant interest in recent years. Efforts in improving security, such as automatic surveillance and the use of biometrics in identification, are partly responsible for this increased interest. However, several challenges remain in improving the accuracy of face recognition under illumination changes, variations in pose, occlusions (including self-occlusion), and image resolution. Many face recognition algorithms have been developed and each has its strengths and weaknesses.

It is well-known that the face exhibits reflectional symmetry about a bilateral symmetry axis. Several authors have noted the role of symmetry in nature [6, 7] and particularly

in human face attractiveness [14, 21, 23]. Additionally, research into using the symmetry of the face to assist with face detection and face recognition has been previously studied. Chen *et al.* [3] developed a method to automatically compute the bilateral symmetry axis (or plane for 3D data), which is particularly useful in this research. The symmetry of the face is used by Zhao and Chellappa [27] to detect and remove illumination effects in faces to improve recognition with what they term Symmetric Shape-from-Shaping. The use of symmetry has also proven useful to extract the facial profile for face recognition, such as in [19, 26]. In [20], Ramanathan *et al.* perform similarity measures between images of the same individual to study the affects of age, disguise, illumination and pose on the face. Our work is different in that we seek to study the symmetry of the face between subjects and the overall impact of symmetry on face recognition. The inspiration of this paper is mostly from the average-half-face, which is based on our previous work [11, 12, 13].

The average-half-face exploits facial symmetry by dividing the frontal full face into two halves about the bilateral symmetry axis, mirroring one of the halves across the symmetry axis, and then averaging the two resulting images. The use of the average-half-face in face recognition research has shown a potential increase in accuracy and decrease in storage and computation time as compared to using the original full face. However, it is not clear why the average-half-face would provide an advantage. Therefore, a statistical analysis of the symmetry of the face is needed.

In this paper we will introduce new measures, or scores, for symmetry of the face to assist in this analysis. The scores are computed using only the pixel values (and the weighted average of pixel values) of the images in the database. A 3D face database is used in our purposes to remove unwanted errors in the symmetry calculation from issues present in 2D images, such as illumination. Using these scores, we perform statistical tests to compare the symmetry of the face in several subgroups of the database and then compare the face recognition results on the same subgroups. A face symmetry analysis of this type has never

been attempted as far as we know. We show a significant difference in face symmetry scores between the men and women subgroups and compare the face recognition results. We then divide the database into the most symmetric subjects and least symmetric subjects according to the symmetry scores and compare their face recognition results. We find that using symmetry in face recognition, by utilizing the average-half-face, is universally beneficial in our experiments.

This paper is organized as follows. First, we will briefly introduce the face recognition problem and results that have motivated this work. This includes the average-half-face, the face recognition algorithms used in our work and the database used. The measures of symmetry will then be introduced for use in our statistical analysis of the symmetry of the face between subgroups. Then we compare the face recognition results of the subgroups and how they relate to the symmetry scores. We end with a discussion of the results and conclude with future work on this topic.

2. Face Recognition

In this section, we will motivate the problem of measuring and analyzing the symmetry of the face by examining recent results from a face recognition algorithm that utilizes face symmetry. First we will describe one use of the symmetry of the face in face recognition. Then we will briefly describe the face recognition algorithm used for comparison and the face database used in our experiments.

2.1. Average-Half-Face

The average-half-face is an average of the two halves of the face along a bilateral symmetry axis of the face, which has proven to produce better results than the original full face for 3D face data, as shown in Figure 2. We will briefly describe the pre-processing algorithm that is used to form the average-half-face. More examples and details can be found in [11, 12].

Algorithm Outline:

1. Center and orient the face in the image. This is done in our work by using the location of the tip of the nose.
2. Partition the image into two equal left and right halves, I_l and I_r .
3. Reverse of the ordering of (or mirror) the columns of one of the images (in this case the right image), producing I_r' .
4. Average the resulting mirror right half image (I_r') with the left half image (I_l) to produce the average-half-face.




Figure 1. (a) 3D example full face image; (b) its average-half-face




Figure 2. Accuracy of Full Face and Average-Half-Face on 3D Face database

The average-half-face is then used in place of the full face image for face recognition. Figure 1 displays an example original full face image from the 3D database and the resulting average-half-face using the steps outlined above.

2.2. Face Recognition Algorithms

In our experiments eigenfaces [24] is used as the face recognition algorithm. We chose to use eigenfaces because of its simplicity and because it highlights the difference in face recognition accuracy between using the full face and the average-half-face, as shown in Figure 2. Eigenfaces is based on principal components analysis (PCA) and is a common face recognition algorithm for benchmarking. PCA captures as much of the variance in the data as possible in as few principal components as possible. All images are projected into the same subspace defined by a chosen number of principal components and test images are classified in the subspace using nearest neighbors.

The face recognition algorithm used in these experi-

ments uses eigenfaces to form two subspaces from the full face and average-half-face. Then, the classification of test faces is done separately in these subspaces. Each of the remaining five face recognition methods that were used to generate the face recognition results in Figure 2 will be briefly described below for completeness.

Multilinear PCA (MPCA) [16] and MPCA-LDA extend PCA to use the entire face image as a tensor object (or multilinear arrays) in order to try to preserve the relationship between neighbor pixels. Fisherfaces [2], another popular method in the literature used for benchmarking, is based on linear discriminant analysis (LDA) and thus attempts to utilize class information in the data to maximize between class scatter while minimizing within class scatter. Independent Components Analysis (ICA) [1] models each face image as a linear combination of non-Gaussian random vectors where the weights of the linear combinations of the training and testing images are used for identification. In each of the above subspace face recognition methods, nearest neighbors was used as the classifier.

2.3. Database

In our experiments we have utilized a 3D face range image database known as the “Texas 3D Face Recognition Database” [8, 9, 10]. A 3D range image is an image in which each pixel represents the depth from the camera. The database consists of a total of 1126 images of 104 subjects. There are anywhere from 1 to 55 images per subject. For the face recognition task, we trained the algorithms using a combination of 360 images from 12 randomly chosen subjects. The gallery is formed from a single neutral expression from each of the subjects and the probes consist of the remaining images.

The images from the 3D database were preprocessed to be centered in the image using the tip of the nose location and an oval mask was applied to remove background noise. Additionally, preprocessing was performed on the images to remove any asymmetric noise that was generated during the scanning process, such as that shown in Figure 3 (a). An example image resulting in the above preprocessing is shown in Figure 3 (b).

In Figure 2, we can clearly see that the average-half-face outperforms the full face in every method used with the 3D face database. More information concerning these results and how they were obtained can be found in [11]. Hopefully learning the statistical differences between subjects within this database will help to uncover the difference in face recognition accuracy when using the average-half-face for recognition.

3. Measuring Symmetry

Defining a measure for the symmetry of an object has been previously investigated. However, it is difficult to find




Figure 3. (a) Example asymmetric face; (b) Preprocessed face

a measure that encapsulates the symmetry of the face in a single number that is used to easily compare the symmetry between different faces. Authors have defined symmetry measures [18] that are useful for finding the symmetry of a single object in an image, but are not useful in comparing symmetry across images and objects, or they are not easily adapted for use with faces [5]. Some symmetry measures are based on feature points on the face and the relationships between these points, such as in [22]. However, these feature points are not readily available on every face database and require manual supervision for reliable accuracy. We have adapted one previous method for measuring symmetry, known as the density difference (D), from the authors in [17]. The measure is formulated by

$$D(i, j) = I(i, j) - I'(i, j), \quad (1)$$

where $I(i, j)$ is a pixel from one half of the image and $I'(i, j)$ is a pixel from the mirror of the other half of the image. The resulting density difference D is itself an image which displays the asymmetry present in the face. However, we desire a single value, or score, for the symmetry of each individual face for comparing the symmetry of many faces. Therefore, we define the following scores:

- Sum of absolute differences (s-score)
- Symmetry proportion (p-score)
- s-score applied to Gaussian smoothed image (sg-score)
- p-score applied to Gaussian smoothed image (pg-score)

The s-score is a simple extension of the density difference and is defined as:

$$s = \sum_{i,j} |D(i, j)|. \quad (2)$$

In addition to the s-score, we introduce a symmetry proportion score (p-score) that is bounded between 0 and 1 and may give a better intuition for the symmetry of the face. The p-score is defined as

$$p = 1 - \frac{\sum_{i,j} T(i,j)}{N}, \quad (3)$$

where $T(i,j)$ is 0 if the absolute difference of the pixels is less than a certain threshold and 1 otherwise and N is the total number of pixels used in the symmetry score. In the experiments used in this paper, the threshold chosen was 10. From this definition, it is apparent that faces that are highly symmetric will give a p-score that is close to 1.

Because these two measures are pixel based and therefore can be sensitive to noise in the image, we also apply them to a Gaussian smoothed image with a window size of 7 pixels and a sigma of 7 pixels. In the remaining sections of the paper, the results of the scores on these smoothed images are reported as the sg-score and pg-score, respectively. These two scores may be less sensitive to errors from image alignment and the scanning process and can be thought of as comparing 7x7 patches of each side of the face. Note that the original s-score and p-score are essentially the same as the sg-score and pg-score, respectively, if the Gaussian smoothing filter has a window size of 1 pixel.

4. Statistical Analysis

As previously discussed, we wish to perform statistical analysis on the symmetry scores that we obtain from face images. For this work, we test several subgroups as follows. First we would like to test if there is a significant difference in the symmetry between men and women in the database. Second, we will use the symmetry scores of each of the subjects to partition the database (including all subjects) into the most symmetric subjects and the least symmetric subjects. This is done by using the average s-score, p-score, sg-score and pg-score of each subject and then sorting the subjects based on these average scores. The most and least symmetric subjects based on using all four symmetry scores are displayed in Figure 4. We will first perform tests for normality on each of the groups and then perform the appropriate hypothesis tests.

4.1. Tests for Normality

We suspect that the samples of the symmetry scores will not be normally distributed because of the upper and lower bounds on the scores. However, it is necessary to establish whether the samples are normally distributed so that the correct hypothesis test is chosen.

The first test for normality is to plot the histograms of the samples for visual inspection. Figures 5, 6, 7 and 8 display the histograms of the s-, p-, sg- and pg-scores, respectively,




Figure 4. (a) Most symmetric and (b) least symmetric subject from the database according to symmetry scores




Figure 5. Histogram of s-score from Men and Women Images

for the samples of men and women from the database. From the figures, it is clear that the distributions are not likely to be normal. However, this is difficult to tell from inspection alone, so we employ two statistical methods to test for normality.

The common Kolmogorov-Smirnov test [15] is first used to test normality. In all subgroups the null hypothesis that each subgroup (individually tested) was sampled from a normal distribution was rejected with an $\alpha = 0.05$. Several authors, however, have noted issues with the Kolmogorov-Smirnov test, so we have also used the recommended D’Agostino-Pearson normality test [4]. The results of this test were the same as that of the Kolmogorov-Smirnov test, so we conclude that our data is not sampled from a normal distribution. Therefore, we must take care in choosing the appropriate statistical test to discover if there is a significant difference between the subgroups.

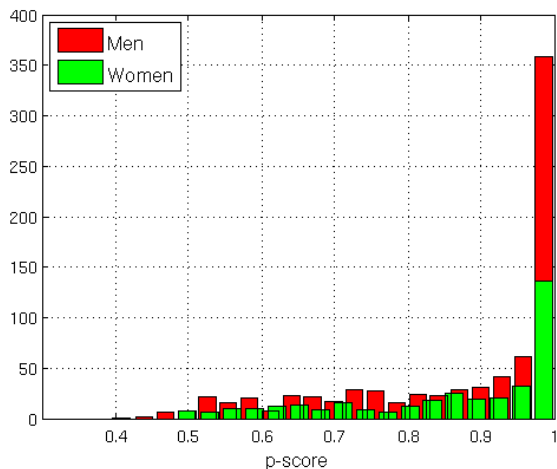


Figure 6. Histogram of p-score from Men and Women Images

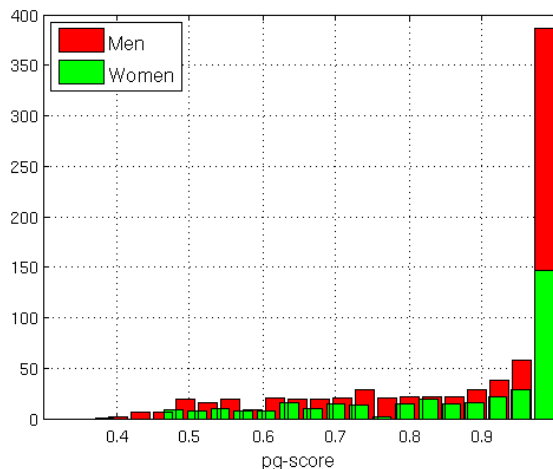


Figure 8. Histogram of pg-score from Men and Women Images




Figure 7. Histogram of sg-score from Men and Women Images

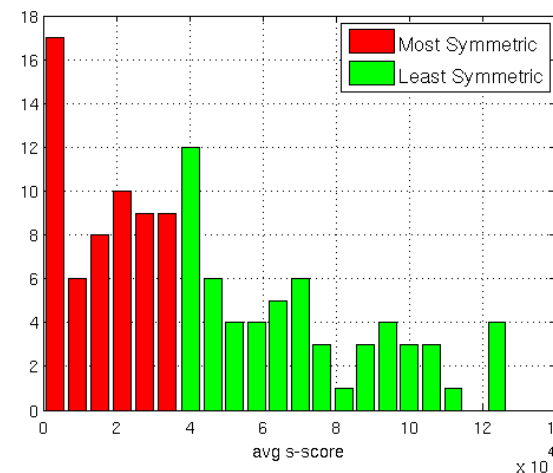


Figure 9. Histogram of s-score from Most and Least Symmetric Subjects

4.2. Paired Two Sample Hypothesis Test

As indicated in the previous section, each of the subgroups that we wish to compare have been determined to not have been sampled from a normal distribution. Therefore, we cannot use statistical tests such as Student's t-test to discover if there is a significant difference between two subgroups. We have chosen to use the nonparametric Wilcoxon test [25] to test the null hypothesis that two populations have the same continuous distribution. We have utilized a Matlab[®] implementation of this test to produce our results.

The results have been summarized in Table 1. The null hypothesis in each of the tests is that the samples from each of the subgroups is drawn from the same distribution with $\alpha = 0.05$. A value of $h = 0$ means that we cannot reject the null hypothesis. A value of $h = 1$ tells us to reject the null

hypothesis and conclude that the two subgroups are drawn from two significantly different distributions. The p-value all of the hypothesis tests are also included for the reader's reference.

In order to better understand this result and why a test of this type is needed, we have included the means and medians of each of the distributions in Tables 2 and 3. In a comparison such as that of men and women in the 3D database using the p-score, the means and medians seem quite close. However, as observed from Table 1, the subgroups are significantly different.

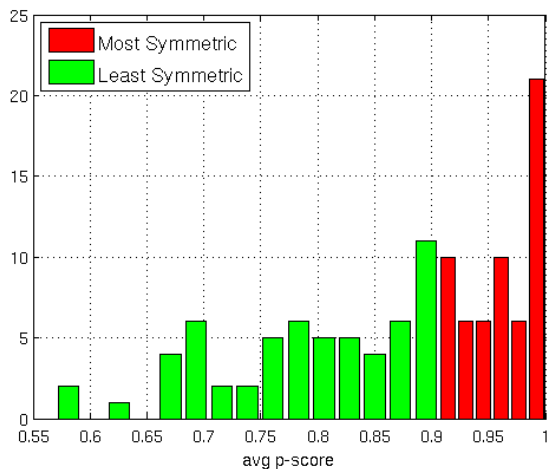


Figure 10. Histogram of p-score from Most and Least Symmetric Subjects




Figure 12. Histogram of pg-score from Most and Least Symmetric Subjects




Figure 11. Histogram of sg-score from Most and Least Symmetric Subjects

Table 1. Wilcoxon Test Results

score	Men v. Women		MostSym v. LeastSym	
	h	p-val	h	p-val
s	1	0.0097	1	<0.0001
p	1	0.0257	1	<0.0001
sg	1	0.0116	1	<0.0001
pg	1	0.0206	1	<0.0001

4.3. Face Recognition Results

We have shown that the subgroups “men” and “women” within the 3D database are sampled from two statistically different distributions. Additionally we have partitioned the database into most symmetric and least symmetric subjects

Table 2. Distribution Means

score	Men v. Women		MostSym v. LeastSym	
	Men	Women	MS	LS
s	4.4E4	5.3E4	1.7E4	7.0E4
p	0.873	0.859	0.963	0.792
sg	4.2E4	5.1E4	1.5E4	6.7E4
pg	0.874	0.856	0.968	0.788

Table 3. Distribution Medians

score	Men v. Women		MostSym v. LeastSym	
	Men	Women	MS	LS
s	2.3E4	3.8E4	1.7E4	6.4E4
p	0.959	0.920	0.959	0.797
sg	2.0E4	3.4E4	1.6E4	6.1E4
pg	0.966	0.929	0.971	0.794

according to the average s-, p-, sg- and pg-score values of each of the subjects. Now we will investigate the face recognition results of each of these subgroups independently and compare their results using both the original full face and the average-half-face.

The face recognition results are obtained in the following way using eigenfaces and nearest neighbors as the classifier. Each of the steps below is repeated separately for the full face and the average-half-face images. First, a common “face space” is formed from a random selection of 12 subjects with 30 images each and all images, including gallery and probe images, are projected into this space for recognition. The gallery is composed of 1 neutral image per subject in the subgroup and the remaining images are used as probes. Therefore, low recognition rates are expected since each subject only has a single gallery image. Nearest neighbors is used as the classification method.

Table 4. Face Recognition Accuracy on Subgroups

Subgroup	FF	AHF	p-value
Men	41.7	60.5	<0.0001
Women	38.2	42.5	0.4336
MostSym-s	35.0	56.3	<0.0001
LeastSym-s	46.0	54.2	0.0264
MostSym-p	42.3	60.0	<0.0001
LeastSym-p	46.4	60.9	0.0023
MostSym-sg	34.3	56.3	<0.0001
LeastSym-sg	45.5	54.0	0.0218
MostSym-pg	33.0	58.2	<0.0001
LeastSym-pg	49.9	58.0	0.0393

Table 5. P-values for Face Recognition Significance Between Most & Least Symmetric Subgroups

score	FF	AHF
s	0.0053	0.6614
p	0.3591	0.8650
sg	<0.0001	0.9976
pg	<0.0001	0.9817

The results from the face recognition on the subgroups using both the original full face (FF) and average-half-face (AHF) are shown in Table 4, where the most symmetric and least symmetric subgroups are labeled as “MostSym- ” and “LeastSym- ” followed by the type of score used for the partitioning of the subjects, i.e., “s”, “p”, “sg” and “pg”. The fourth column of Table 4 displays the p-values from the two proportion hypothesis tests which tests whether the differences in accuracy between using the full face and the average-half-face are statistically significant. The second column of Table 5 displays the p-values from the hypothesis test between the most and least symmetric face subgroups using the full face, while the third column displays the results when the average-half-face is used. From the results, it is clear that the average-half-face outperforms the full face for every subgroup involved.

5. Discussion

From the paired two sample hypothesis tests performed on the men and women subgroups, as well as the most symmetric and least symmetric subgroups, using the s-score, p-score, sg-score and pg-score, it is clear that each of the subgroups are sampled from statistically different distributions, as shown in Table 1. Therefore, one conclusion that can be drawn is that the measures of symmetry are consistent. Another is that, for this particular database, the images of men are more symmetric than that of the women. When comparing the face recognition results in Table 4 between men and women, a similar difference is noted. When using the full face for recognition, the men have a slightly larger recognition rate than the women, though not statistically significant

(p-value of 0.4419). The surprising result is that when using the symmetry of the face with the average-half-face for recognition, the men have a significantly higher recognition rate of 60.5% compared with the women’s recognition rate of 42.5% (with a p-value < 0.0001). When comparing the accuracy of using the full face with the average half face for men and women subgroups, only the result for men is statistically significant (p-value < 0.0001).

Partitioning the database into most symmetric and least symmetric subgroups with the symmetry scores produces clearly different distributions as shown in Table 1 and in Figures 9, 10, 11 and 12. When using these subgroups for face recognition, another surprising result is discovered. The face recognition results are higher for the least symmetric faces than that of the most symmetric faces when using the full face. This is potentially explained by thinking of a simple example. In the full face images, features that are present on only one half of the face, such as a mole or scar, are more discriminant than features that are shared on both halves of the face, so we might expect face recognition algorithms to perform better on faces which are more asymmetrical. However, the results when using the average-half-face are basically the same between the most symmetric and least symmetric halves. Of course, in the case of both symmetry scores, using the average-half-face is far more beneficial than using the full face. So, it appears that the average-half-face is not biased towards the symmetry score of the subject when performing face recognition and additionally provides a boost in accuracy to both the most symmetric and least symmetric subgroups. As shown in Table 4, the difference between the full face and average-half-face accuracies are statistically significant with $\alpha = 0.05$, except for the women subgroup. From Table 5, the difference between face recognition accuracies of the most and least symmetric subgroups are significant only for the full face results.

Therefore, when performing a face recognition task such as that described in this paper, performing a face symmetry analysis on the faces in the database may help predict the face recognition performance when using a frontal full face image. However, the average-half-face appears to result in higher face recognition accuracy regardless of the symmetry inherent in the database.

6. Conclusions and Future Work

We have presented in this paper a statistical analysis of the relationship between the symmetry of the face and face recognition. We have introduced new symmetry scores and used them to compare men and women subgroups as well as most symmetric and least symmetric subgroups. We have found a statistical significance between the face symmetry of men and women subjects in the 3D database as well as differences in face recognition accuracy. The least symmetric subjects produce higher face recognition accu-

racy than the most symmetric subjects when using the full face. However, face recognition accuracy is universally improved when utilizing the average-half-face in our experiments over the full face. Future work will include introducing additional symmetry measures as well as extending this analysis to more databases of 2D and 3D faces. Also, even though these 3D images were taken in a very controlled environment, additional normalization of pose may be necessary before computing the symmetry score to account for rotational errors or mislabeled fiducial points on the face, as shown in Figure 4 (b). The ultimate goal of this work would be to create a correlation between the symmetry of the face and face recognition that could be used to improve the overall face recognition accuracy.

References

- [1] M. S. Bartlett, J. R. Movellan, and T. J. Sejnowski. Face recognition by independent component analysis. *IEEE Transactions on Neural Networks*, 13:1450–1464, 2002. 75
- [2] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. In *ECCV '96: Proceedings of the 4th European Conference on Computer Vision-Volume I*, pages 45–58, London, UK, 1996. Springer-Verlag. 75
- [3] X. Chen, P. J. Flynn, and K. W. Bowyer. Fully Automated Facial Symmetry Axis Detection in Frontal Color Images. In *Automatic Identification Advanced Technologies, 2005. Fourth IEEE Workshop on*, pages 106–111. IEEE, 2005. 73
- [4] R. D’Agostino. Tests for normal distribution. *Goodness-Of-Fit Techniques*, pages 367–419, 1986. 76
- [5] S. Dakin and A. Herbert. The spatial region of integration for visual symmetry detection. *Proceedings of the Royal Society B: Biological Sciences*, 265(1397):659, 1998. 75
- [6] T. Davis and C. Ramanujacharyulu. Statistical analysis of bilateral symmetry in plant organs. *Sankhyā: The Indian Journal of Statistics, Series B*, 33(3):259–290, 1971. 73
- [7] P. Endress. Symmetry in flowers: diversity and evolution. *International Journal of Plant Sciences*, 160(6):3–23, 1999. 73
- [8] S. Gupta, J. K. Aggarwal, M. K. Markey, and A. C. Bovik. 3D face recognition founded on the structural diversity of human faces. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:1–7, 2007. 75
- [9] S. Gupta, K. R. Castleman, M. K. Markey, and A. C. Bovik. *Texas 3D Face Recognition Database*, 2010. <http://live.ece.utexas.edu/research/texas3dfr/index.htm>. 75
- [10] S. Gupta, K. R. Castleman, M. K. Markey, and A. C. Bovik. Texas 3D Face Recognition Database. *IEEE Southwest Symposium on Image Analysis and Interpretation*, pages 97–100, May 2010. 75
- [11] J. Harguess and J. K. Aggarwal. A case for the average-half-face in 2D and 3D for face recognition. *IEEE Computer Society Workshop on Biometrics (in conjunction with CVPR)*, 2009. 73, 74, 75
- [12] J. Harguess and J. K. Aggarwal. The Average-Half-Face in 2D and 3D Face Recognition. In *Pattern Recognition and Machine Vision*, pages 135–148. River Publishers, 2010. 73, 74
- [13] J. Harguess, S. Gupta, and J. K. Aggarwal. 3D face recognition with the average-half-face. *International Conference on Pattern Recognition ICPR*, 2008. 73
- [14] J. Honekopp, T. Bartholome, and G. Jansen. Facial attractiveness, symmetry, and physical fitness in young women. *Human Nature*, 15(2):147–167, 2004. 73
- [15] H. Lilliefors. On the Kolmogorov-Smirnov test for normality with mean and variance unknown. *Journal of the American Statistical Association*, 62(318):399–402, 1967. 76
- [16] H. Lu, K. N. Plataniotis, and A. N. Venetsanopoulos. MPCA: Multilinear Principal Component Analysis of Tensor Objects. *IEEE Trans. on Neural Networks*, 19(1):18–39, 2008. 75
- [17] S. Mitra, N. Lazar, and Y. Liu. Understanding the role of facial asymmetry in human face identification. *Statistics and Computing*, 17(1):57–70, 2007. 75
- [18] D. O’Mara and R. Owens. Measuring bilateral symmetry in digital images. In *TENCON’96. Proceedings. 1996 IEEE TENCON. Digital Signal Processing Applications*, volume 1, pages 151–156. IEEE, 2002. 75
- [19] G. Pan and Z. Wu. 3D face recognition from range data. *Int. J. Image Graphics*, 5(3):573–594, 2005. 73
- [20] N. Ramanathan. Facial similarity across age, disguise, illumination and pose. In *Proceedings of International Conference on Image Processing*, 1999. 73
- [21] G. Rhodes, F. Proffitt, J. Grady, and A. Sumich. Facial symmetry and the perception of beauty. *Psychonomic Bulletin and Review*, 5:659–669, 1998. 73
- [22] K. Schmid, D. Marx, and A. Samal. Computation of a face attractiveness index based on neoclassical canons, symmetry, and golden ratios. *Pattern Recognition*, 41(8):2710–2717, 2008. 75
- [23] R. Thornhill and S. Gangestad. Facial attractiveness. *Trends in Cognitive Sciences*, 3(12):452–460, 1999. 73
- [24] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991. 74
- [25] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945. 77
- [26] L. Zhang, A. Razdan, G. Farin, J. Femiani, M. Bae, and C. Lockwood. 3D face authentication and recognition based on bilateral symmetry analysis. *The Visual Computer*, 22(1):43–55, 2006. 73
- [27] W. Zhao and R. Chellappa. Illumination-insensitive face recognition using symmetric shape-from-shading. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 1:1286, 2000. 73