

A Quantified Study of Facial Asymmetry in 3D Faces

Yanxi Liu and Jeff Palmer

CMU-RI-TR-03-21 2

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

©2003 Carnegie Mellon University

This research is supported in part by an ONR research grant N00014-00-1-0915 (HumanID), and in part by an NSF research grant IIS-0099597.

University of CS
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

With the rapid development of 3D imaging technology, the wide usage of 3D surface information for research and applications is becoming a convenient reality. This study is focused on a quantified analysis of facial asymmetry of more than 100 3D human faces (individuals). We investigate whether facial asymmetry differs statistically significantly from a bilateral symmetry assumption, and the role of global and local facial asymmetry for gender discrimination.

1 Introduction

The recent advent of 3-dimensional imaging technology has spurred on several new research questions regarding the development of precise human detection and identification algorithms. Most of the object detection research done in 3D has thus far focused on non-human objects. Several 3D descriptors have been developed for classification of objects that make use of localized surface matching, e.g. [7], and global symmetrical attributes of a surface [8]. Little research has been done in the area of human face identification using 3D media. However, 3D imaging technology has advanced so rapidly that performing a quick 3D scan of every person who passes through an airport security gate may become routine in the near future. Some researchers have already shown certain success in 3D face comparison [6, 1].

One of the common characteristics of human faces is that they are approximately bilaterally symmetrical. Previous work has explored both human facial symmetry and asymmetry in different applications. Vetter and Poggio [17] suggest that object detection within the 2D framework can be greatly simplified by exploiting the known symmetries of an object in 3D. In [16], Seitz and Dyer assume Mona Lisa's face is bilaterally symmetrical to achieve realistic 3D morphing. Zhao and Chellappa [18] have also taken advantage of facial symmetry in 3D reconstruction of human faces. On the other hand, facial asymmetry has long been reported [4, 2, 3]. Especially, expression lateralization is commonly accepted in the psychology literature, where some findings show facial expression is more intense on the left side. For human identification purposes, half faces have been tried, with some reported differences in recognition rates when left and right face images are used alone [13, 14], while others [5] did not report any differences. More recently, Liu *et al* have demonstrated that quantified facial asymmetry contains discriminating information for human identification under expression variations, and for expression classification across individuals on 2D videos [12, 9, 11, 10]. They show statistically significant improvements can be achieved by combining facial asymmetry with classic human identification algorithms [11, 10].

The diverse applications and opposing assumptions on human facial bilateral symmetry and asymmetry are the primary motivations for this current investigation on 3D human faces. The objectives of this study are to

1. define a set of 3D facial asymmetry measurements that are readily computable directly from 3D digitized data sets;
2. find out whether the computed 3D facial asymmetry differs statistically significantly
 - from a facial bilateral symmetry assumption, and
 - between genders (male versus female)

3. learn correspondences between projected facial asymmetry (2D) and pose estimation in 3D.

The emphasis of this paper is on the first two objectives. Some initial observations on the third are also demonstrated.

2 3D Face Data

The database of human 3D face images that we use in this study is maintained by researchers in the department of Computer Science at the University of South Florida, and sponsored by the U.S. Department of Defense, Defense Advanced Research Projects Agency (DARPA). The database contains one digitized scan (image) for each of $N = 111$ subjects (74 male; 37 female). All subjects have a "neutral" facial expression.

2.1 Data Format

Each image file is stored in the database in the form of a triangle mesh, which consists of x -, y -, z - coordinate triples (observed locations of the surface of the face), and a list of triangles (plane segments) connecting these points. The triangle mesh representation is commonly used as a piecewise linear approximation to the true surface of an object in 3D space. The raw data consists of approximately 200,000 mesh points and 250,000 triangles per subject. The images have been normalized with respect to seven user-selected landmark locations on the face via an affine transformation of the original data [1], such that they are bilaterally located with respect to the YZ-plane.

Refer to Figure 1 for an example of a rendered image from the database, with x -, y -, z -coordinate axes superimposed on a 3D head. The x axis represents left-to-right displacement, y axis front-back, and z axis up-down of the head.

2.2 Data Preprocessing

To increase computational efficiency, we have employed a triangle decimation algorithm [15] to produce a subsample of the original mesh down to 5% of the original data volume.

Since the focus of our analysis in this work is the asymmetrical attributes of the human face, we consider a more conceptually intuitive coordinate system by transforming each 3D image from its Cartesian coordinate system indexed by (x, y, z) to a cylindrical coordinate system (θ, r, z) , i.e.

$$\theta = \tan^{-1} \left(\frac{x}{y} \right), \quad r = \sqrt{x^2 + y^2}, \quad z = z. \quad (1)$$




Figure 1: *Example rendered image from database with x - y - z - coordinate axes superimposed.*




Figure 2: *Left: cropped 3D face of the same subject. Right: meshed 3D face.*

Note, the YZ-plane, where $\theta = 0$, bisects a 3D face.

For convenience of computation, we impose a 2D grid pattern on each 3D face which covers the whole face from chin to forehead in the vertical z direction, and from outside of one eye outer corner to outside the other eye in the θ direction (left of Figure 2). We split the grid into 125 equally spaced intervals in both the θ and z directions, yielding a total of $125 \times 125 = 15,625$ grid points. For each of these 15,625 grid points, we interpolate its r value (distance from the z -axis) using trilinear interpolation on the original surface mesh. See /usr2/figures/Face/3Dface 1 and 2 for a 3D face of the same subject in its rendered, cropped and meshed forms.

If we use F_i to denote the original mesh data and f_i the re-parametrized F_i : each 3D face F_i in the database is now mapped into f_i , a 125 by 125 grid matrix with

$$f_i = r(\theta, z) \tag{2}$$

where θ takes on a discrete set of values in the range $[-0.75, 0.75]$ radians with increment 0.012, and z takes on a discrete set of values in $[-50, 75]$.

3 3D Face Analysis

In this section we define a pair of facial asymmetry measures for 3D faces. These definitions are general and not limited to frontal faces.

3.1 3D Facial Asymmetry Measures

Let γ represent a plane in 3D space passing through the origin. In general, $\gamma(\theta, \phi)$ is a function of rotation about the z -axis (θ) and rotation about y -axis (ϕ). We are currently keeping ϕ fixed at 0 and varying θ only. If we further restrict γ to only take on θ values that coincide with our grid definition, then with respect to each position of the γ plane when $\theta = \theta_\gamma$, each point p_i on a 3D face mesh with coordinates $(\theta_\gamma + \theta_i, r_i, z_i)$ has a corresponding reflected (with respect to γ) point p'_i with $(\theta_\gamma - \theta_i, r'_i, z_i)$, given that the point is within the boundary of the grid.

For a given face f and plane γ , we introduce the following asymmetry measurements:

Height Difference (HD):

$$HD_{f,\gamma}(\theta, z) = |r(\theta, z) - r'_\gamma(\theta, z)|$$

where r'_γ is the r value of face f at $(\theta_\gamma - \theta, z)$, the reflected point of (θ, z) with respect to plane γ . We refer to γ as the *symmetry plane*.

Orientation Difference (OD):

$$OD_{f,\gamma}(\theta, z) = \cos(\beta_{\vec{v}_{\theta,z}, \vec{v}'_{\theta,z}})$$

where $\beta_{\vec{v}_{\theta,z}, \vec{v}'_{\theta,z}}$ is the angle between the normal vector of the face mesh at grid point (θ, z) , and the normal vector of its corresponding reflected point with respect to γ .

For each 3D face f and plane of symmetry γ pair, the $HD(f_i)$ and $OD(f_i)$ functions both yield feature vectors with dimensions as high as the number of grid points. Given the symmetry of the HD, OD definitions the left or right half of HD -face (OD -face) already contains all of the facial asymmetry information in HD -face (OD -face). We refer to these measurements as HD -Face and OD -Face. Figure 3 shows the original mesh, the HD -Face, and the OD -Face for six subjects from the database. The top three rows in the figure are females and the bottom three rows are males.

3.2 Global Facial Asymmetry

To obtain an overall measure of asymmetry for a particular subject with respect to a plane of symmetry, we can take the average of all the values over the grid of HD -Face or OD -Face.

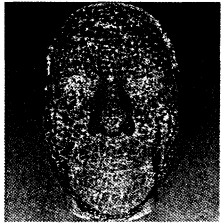

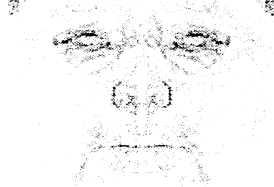
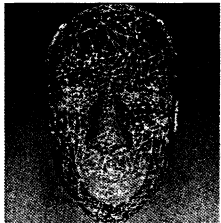


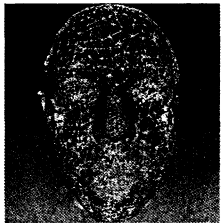
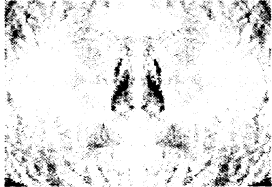

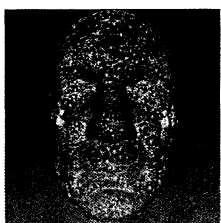
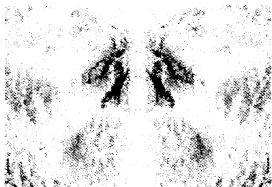

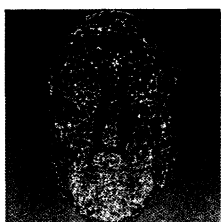
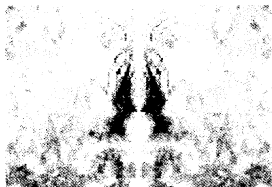
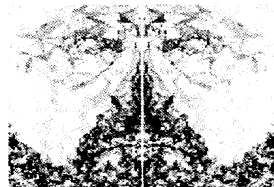
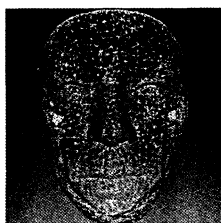


Subject	Original mesh	HD-Face	OD-Face
001, Female			
008, Female			
053, Female			
002, Male			
023, Male			
037, Male			
Subject	Original mesh	HD-Face	OD-Face

Figure 3: Original mesh images, HD-Faces, and OD-Faces for six subjects in the database. Top three rows

The resultant value will be a norm representing $\|f - f_\gamma\|$ where f_γ is the reflection of f through the plane of symmetry γ . In the *HD*-Face case, this norm is the 1-norm. So, we define the following global metrics:

Overall Height Difference (OHD):

$$OHD_{f,\gamma} = \frac{1}{I} \sum_{\theta,z} |r(\theta, z) - r'_\gamma(\theta, z)|$$

Overall Orientation Difference (OOD):

$$OOD_{f,\gamma} = \frac{1}{I} \sum_{\theta,z} \cos(\beta_{\vec{v}_{\theta,z}, \vec{v}'_{\theta,z}})$$

where I is the total number of grid points that have a corresponding reflected point on the grid. For $\theta_\gamma = 0$ (symmetry plane that bisects the head through the midline of the face into a right and left hemisphere), each point on the grid has a corresponding reflected point that is also on the grid (so, $I = 125 \times 62 = 7750$). For values of γ other than 0, however, there will be fewer grid points with which to use in the computation of *OHD* and *OOD*.

4 Analysis Results

We start with an analysis of global 3D facial asymmetry of the *OHD* and *OOD* measurements, followed by a local analysis. We are looking into the question of whether there are statistically significant differences between the measured facial asymmetry of all subjects and the human facial bilateral symmetry assumption. We then compare male and female facial asymmetry globally and locally.

4.1 Global Asymmetry

For each subject in the database with a 3D face f_i , we vary the symmetry plane such that $\theta_\gamma \in [-0.36, 0.36]$ radians. Table 1 shows the means and standard errors of *OHD* and *OOD* for a range of values of θ_γ . The question is: Do the mean asymmetry values differ significantly from the bilateral symmetry assumption?

Table 1: Means and standard errors of *OHD* and *OOD* for varying values of θ_γ .

θ_γ (radians)	OHD		OOD	
	Mean	SE	Mean	SE
-0.363	12.33	2.28	0.73	0.07
-0.302	13.13	2.22	0.67	0.04
-0.242	11.69	2.09	0.64	0.04
-0.181	10.07	1.74	0.58	0.03
-0.121	8.12	1.25	0.54	0.04
-0.060	4.81	0.82	0.41	0.04
0.000	1.44	0.35	0.20	0.03
0.060	4.94	0.80	0.40	0.04
0.121	8.28	1.22	0.54	0.04
0.181	10.29	1.73	0.58	0.03
0.242	11.97	2.07	0.64	0.04
0.302	13.45	2.20	0.67	0.04
0.363	12.75	2.33	0.73	0.07

As expected, with the increase of θ_γ the values of *OHD* and *OOD* increase. It is interesting to note that the *OOD* measurement has a nearly constant variance for different values of γ , while the variance of *OHD* is fanning out as the symmetry plane moves away from the frontal view. Figure 4 shows the *OHD* and *OOD* plots, respectively, for all subjects in the database. Note the single aberrant profile in the *OOD* panel of Figure 4, who upon inspection of Figure 3, appears to have a beard that contributes a large amount of asymmetry in

terms of the vectors normal to the mesh. We have removed this subject from subsequent analysis.

The vertical reference line in Figure 4 signifies the frontal view plane of symmetry, note both the means and the variances of the 3D facial asymmetry measures are non-zero. Figure 5 shows the facial asymmetry distribution histograms with density estimates overlaid on the graphs, when the symmetry plane is the YZ-plane (frontal view). The outlier in the *OOD* measurements is the same subject 023 as in Figure 4.

In terms of frontal faces ($\theta_\gamma = 0$), a perfectly symmetric face would yield a global asymmetry measurement of zero. To assess the statistical significance of the level of overall asymmetry we have observed on the frontal faces, we compare the distribution of the observed values to a null hypothesis of a facial asymmetry distribution with expected value of zero. The level of variation associated with this null distribution is unknown, but we can estimate it via a bootstrap approach by using the variance of the observed values from *OHD* and *OOD* respectively.

The observed values are approximately normally distributed with a lower bound of zero, so we assume that they follow a truncated normal distribution. A test of significance yields a *p*-value of < 0.001 for both *OHD* and *OOD* measurements, suggesting that the frontal faces are **nominally bilaterally asymmetric**.

4.2 Global Frontal Face Asymmetry by Gender

Previous results have reported a facial asymmetry difference between genders (e.g. [4]). In this section we focus our attention on the relationship between facial asymmetry and the gender differences. Once again, we restrict our attention to the global asymmetry measures obtained for the frontal face (plane of symmetry corresponding to $\theta_\gamma = 0$). Figure 6 presents box-plots portraying the distributions of the two global measures by gender.

A simple t-test comparing the mean asymmetry value between males and females is statistically insignificant for the *OHD* measure (*p*-value=0.88), yet is **highly statistically significant for the *OOD* measure** (*p*-value < 0.001).

4.3 Local Frontal Face Asymmetry

We now turn our focus to the localized asymmetry vectors *HD*-Face and *OD*-Face. For this discussion, we will again restrict our attention to the frontal face. See the left column in Figure 7 for the average *HD*-Face and *OD*-Face (averaged over all subjects). The darker regions in the images represent areas of higher asymmetry. The images in the right column of Figure 7 show the top 20% most asymmetrical regions measured by *HD*-Face and *OD*-Face.

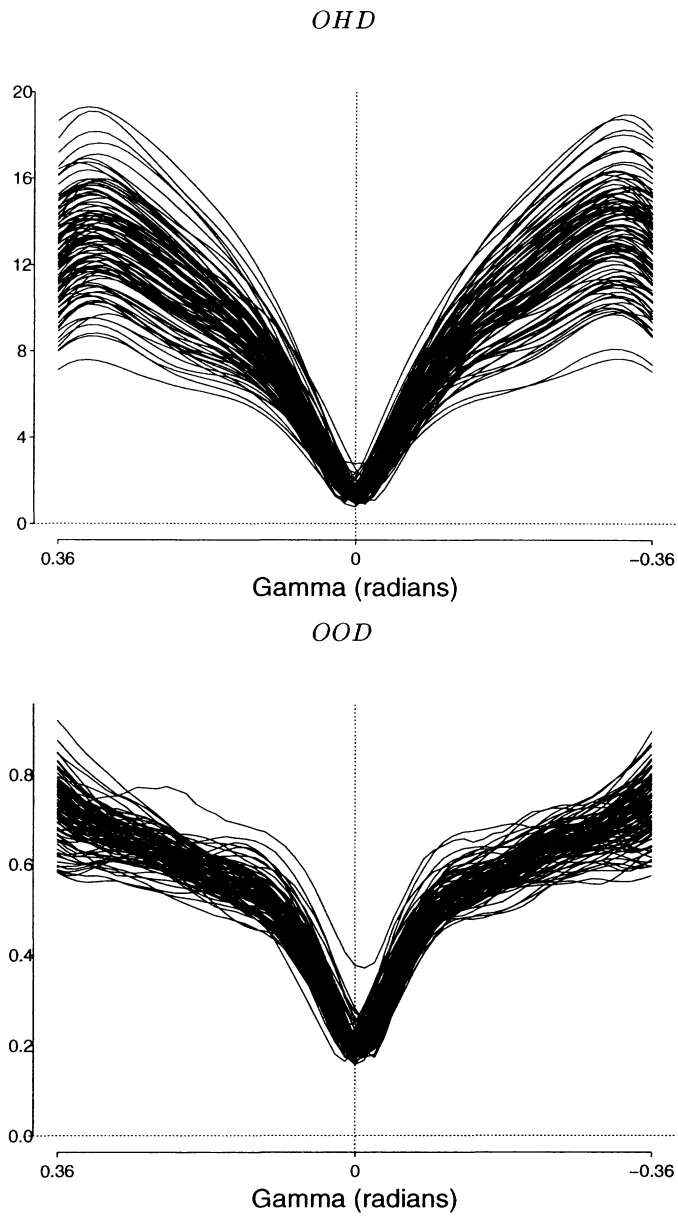


Figure 4: Profile plots of global asymmetry measures. Left panel: OHD, Right panel: OOD.

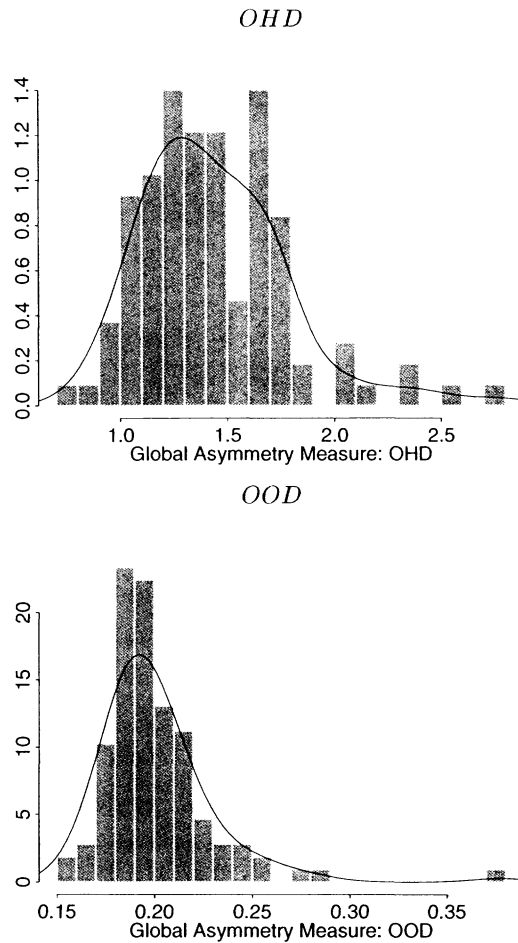


Figure 5: Histograms of global facial asymmetry measures when the symmetry plane $\theta_\gamma = 0$ with density estimates overlaid.

It is clear from these average faces that there are significant trends in the regions of asymmetry for both types of asymmetry measurements. For the *HD-Face*, the regions with the highest asymmetry are around the sides of the nose and in the lower portion of the cheek area. For the *OD-Face*, the regions are the eyes, nostrils, and mouth. The two measurements appear to complement one another quite well in that they capture different information about the regions of asymmetry of the frontal face.

4.4 Local Frontal Face Asymmetry by Gender

We will now investigate the differences in the local asymmetry measures (*HD-Face* and *OD-Face*) between male and female 3D faces. In particular, we are seeking to locate discriminative




Figure 6: *Box-plots for global asymmetry measures divided by gender.*

subspaces of the asymmetry features (facial asymmetry measures belong to certain subregions of the face) where male versus female facial asymmetries differ the most.

Each feature in the full range of *HD-Face* and *OD-Face* vectors is not equally important for discriminating between male and female faces. We attempt to reduce the size of the feature space (dimensionality of *HD-Face* and *OD-Face*) through the means of a variance ratio. In general, for a feature F with C total classes, we define the variance ratio as follows:

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

where $mean_i(F)$ is the mean of feature F 's values in class i . This variance ratio is the ratio of the variance of the feature between classes to the variance of the feature within classes, with an added penalty for features which may have small intra-class variance but which have close inter-class mean values.

Refer to Figure 8 for a feature-by-feature representation of $vr(F)$ applied to the *HD-* and *OD-* Faces with male and female as two gender classes ($C = 2$). The darker regions in the plot signify features that have higher $vr(F)$ values, and thus higher gender discriminating power.

Comparing /usr2/figures/Face/3Dface 8 and 7, one can observe resemblance between the regions on the face with the highest gender discriminating power and the regions of high asymmetry measures.

We may now obtain a feature subspace that is suitable for discriminating between genders by simply looking at those features representing a certain percentage of the largest values of $vr(F)$. If we choose to retain the highest 1% of the features, then we are left with a

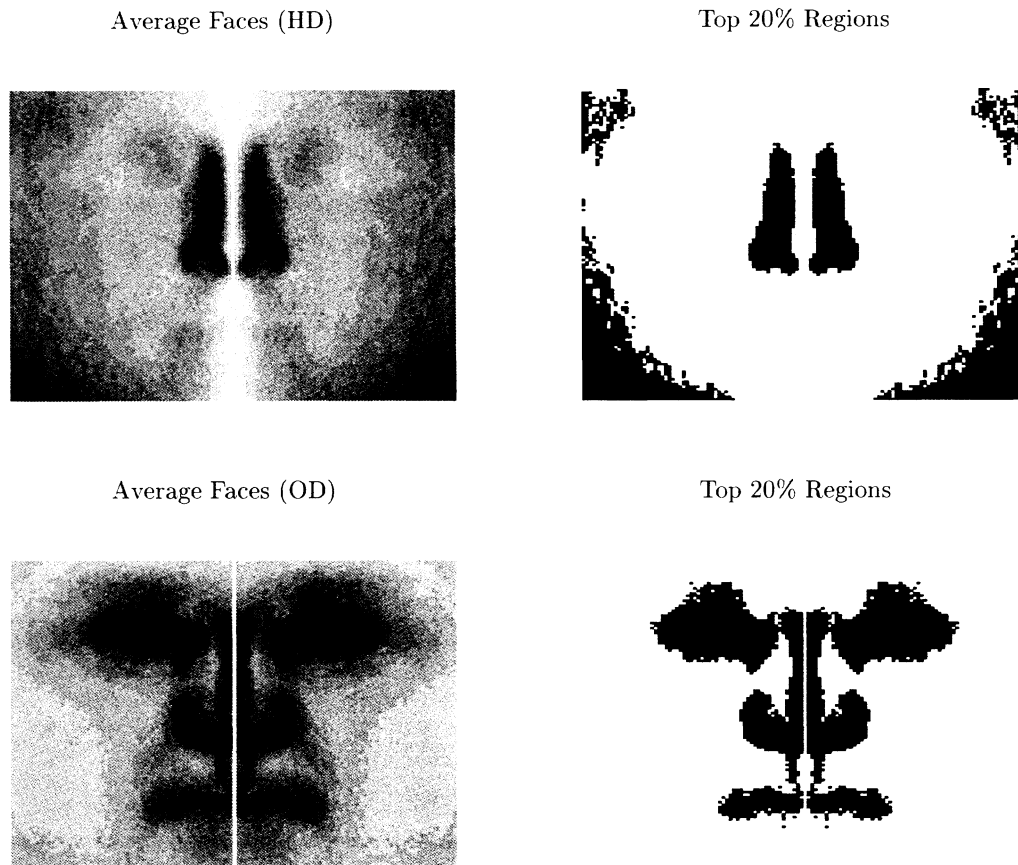



Figure 7: *Average frontal HD-Faces and OD-Faces (left column), and top 20% most asymmetrical regions (right column).*

feature space that has slightly less than 100 dimensions. If we use this subspace in a logistic regression setting with gender acting as the dichotomous outcome, **we observe a highly statistically significant difference between genders (p -value < 0.001 in both cases).** This result suggests that there is enough information in the localized asymmetry measures for gender classification.

5 Discussion and Future Work

While 3D imaging technology is still in its formative stages of development, many current research efforts are showing promising gains towards the ultimate goal of constructing a




vr(F) of HD-face



vr(F) of OD-face



top 1% of vf(HD)



top 1% of vf(OD)

Figure 8: Variance ratio surfaces of *HD-face* and *OD-face* for discriminating between two classes: male versus female faces.

completely autonomous 3D human identification system. We have attempted in this paper to quantify the level of asymmetry possessed by the 3D human face. We have also addressed the issue of whether there exists significant facial asymmetry in frontal faces and differences in asymmetry between males and females.

In this paper, we defined two types of 3D facial asymmetry measures: height difference faces (*HD-face*) and orientation difference face (*OD-face*). From the statistical analysis, *OD-face* presents itself as a more compact representation for facial asymmetry (Figure 4). Thus it is a good candidate for asymmetry-based pose estimation. *OD-face* also shows more significant departure from the bilateral symmetry assumption and better discriminating power between genders. On the other hand, *HD* face requires less computation cost and is

less sensitive to outliers.

Our result on 3D human faces of more than 100 subjects shows that human facial asymmetry is statistically significant, therefore facial asymmetry should not be ignored without a justification. If properly utilized, it is foreseeable that facial asymmetry may provide non-trivial contributions to human face analysis. One such example has been demonstrated for human identification under expression variations [11].

Our result on 3D facial asymmetry divided by gender also supports early research in psychology that the male face possesses a larger amount of asymmetry than does the female faces. We show that a statistically significant difference arises between male and female faces in terms of the *OOD* measurement, but not so with respect to the *OHD* measure. By analyzing 3D facial asymmetry locally, we are able to identify which part of the face that is most asymmetrical, and which part of the 3D face that is most discriminative in gender classification (Figure 8). The latter finding of regions around the nose bridge echos the discriminative asymmetry measures found from expression videos for human identification [11]. We show that a statistically significant difference exists for both *HD* and *OD* facial asymmetry measurements between males and females when we look at an appropriate subset of the feature space.

Our current work is focusing on the use of facial asymmetry measures for pose estimation in aiding 2D human identification. Some initial attempt along this direction seems to be promising. Figures 9 and 10 shows the pose variation measured by the corresponding asymmetry faces while the asymmetry plane γ rotates in 3D space.

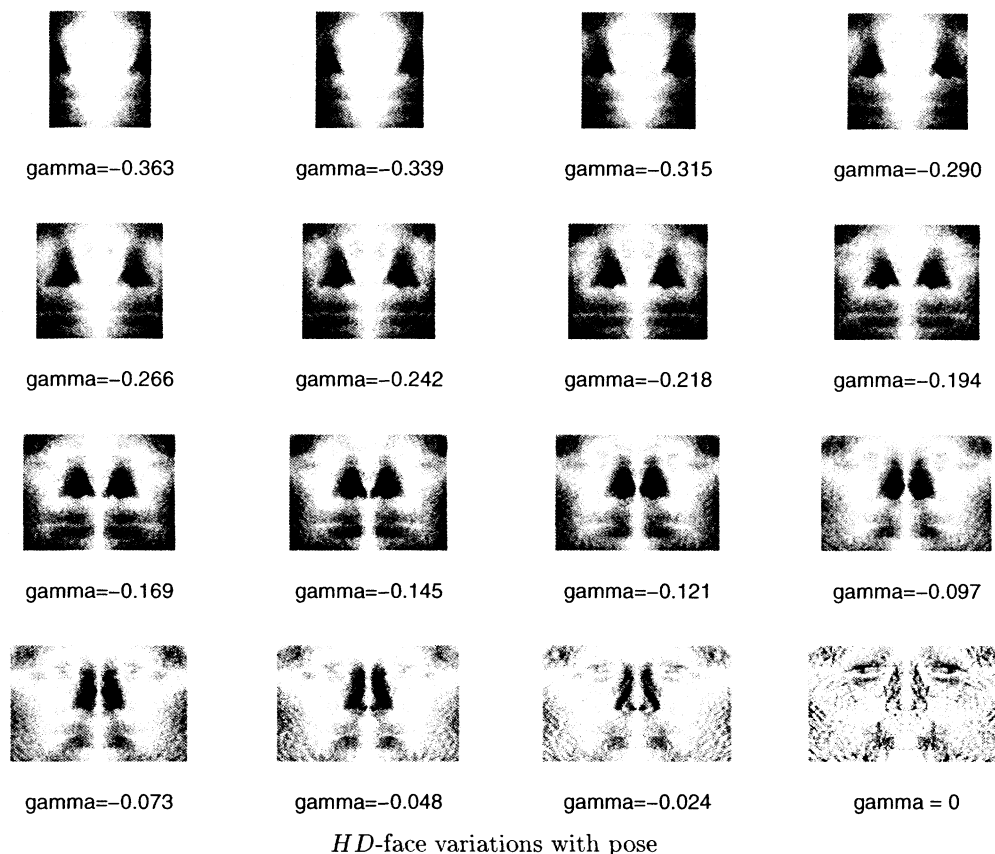


Figure 9: *Pose variations can be determined by the variation of facial asymmetry measures under a varying symmetry plane. Here the symmetry plane γ changes from $\theta_\gamma = 0$ to -0.363 radians (also see Table 1 for facial asymmetry values).*

References

- [1] V. Blanz and A. Vetter. Morphable model for the synthesis of 3d faces. In *SIGGRAPH '99 Conference Proceedings*, pages 187–194, 1999.
- [2] 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.
- [3] R. Campbell. The lateralization of emotion: A critical review. *International Journal of Psychology*, 17:211,219, 1982.

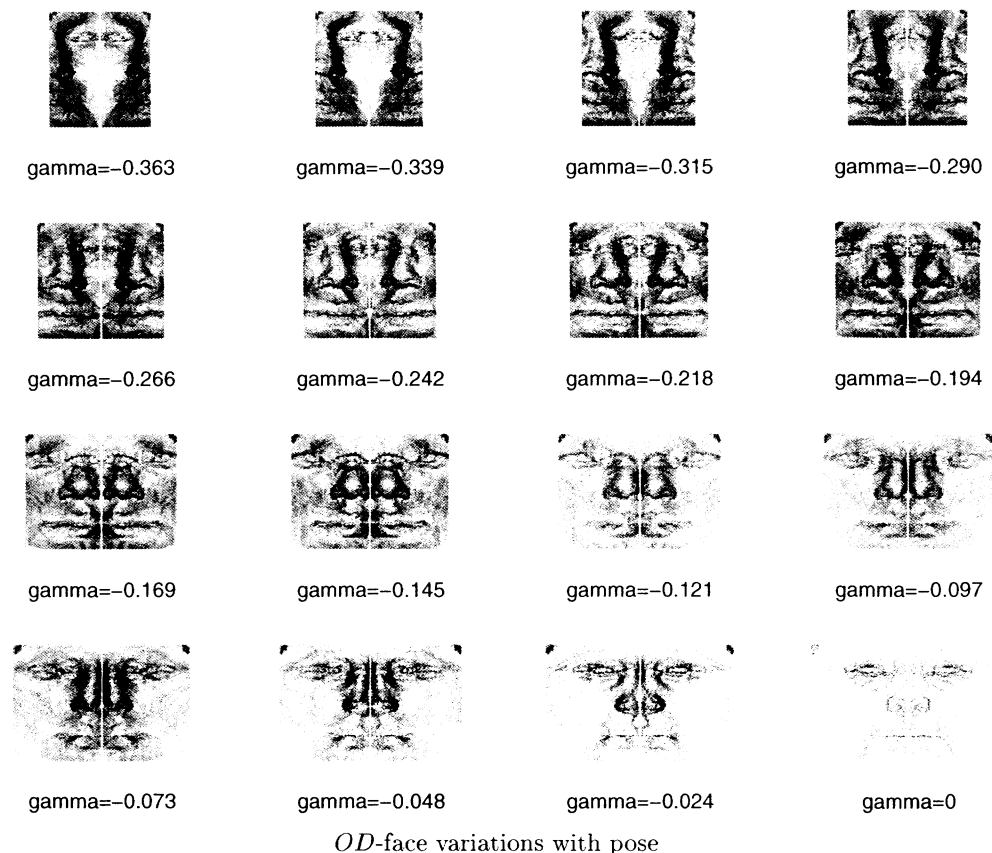


Figure 10: Pose variations can be determined by the variation of facial asymmetry measures under a varying symmetry plane. Here the symmetry plane γ changes from $\theta_\gamma = 0$ to -0.363 radians (also see Table 1 for facial asymmetry values).

- [4] L. G. Farkas. Facial asymmetry in healthy north american caucasians. *Angle Orthodontist*, 51(1):70–77, 1981.
- [5] 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, Waxhington, D.C., May 2002. IEEE Computer Society.
- [6] H. Ip and W. Wong. 3d head models retrieval based on hierarchical facial region similarity. Technical report, City University of Hong Kong., 2000.
- [7] A. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21:433–

449, 1999.

- [8] M. Kazhdan, B. Chazelle, D. Dobkin, A. Finkelstein, and Funkhouser T. A reflective symmetry descriptor. In *Proceedings of Seventh European Conference on Computer Vision*, pages 642–656. Springer, May 2002.
- [9] Y. Liu and S. Mitra. Experiments with quantified facial asymmetry for human identification. Technical Report CMU-RI-TR-02-24, The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 2002.
- [10] Y. Liu and S. Mitra. Human identification versus expression classification via bagging on facial asymmetry. Technical Report CMU-RI-TR-03-08, The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 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*, In press 2003.
- [12] Y. Liu, R.L. Weaver, K. Schmidt, N. Serban, and J. Cohn. Facial asymmetry: A new biometric. Technical Report CMU-RI-TR-01-23, The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 2001.
- [13] A.M. Martinez. Semantic access of frontal face images: The expression-invariant problem. In *Proceedings of IEEE Workshop on Content-based access of images and video libraries*, 2000.
- [14] 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.
- [15] W. Schroeder, J. Zarge, and W. Lorenson. Decimation of triangle meshes. Technical report, General Electric Company., 1992.
- [16] S.M. Seitz and C.R. Dyer. View morphing. *SIGGRAPH*, pages 21–30, 1996.
- [17] T. Vetter and T. Poggio. Symmetric 3d objects are an easy case for 2d object recognition. *Spatial Vision.*, 8(4):443–453, 1994.
- [18] W.Y. Zhao and R. Chellappa. Symmetric shape-from-shading using self-ratio image. *IJCV*, 45(1):55–75, October 2001.