

COMP 766: Shape Analysis in Computer Vision
Term Project

Symmetries of Occluded and Non-Rigid Objects

Mostafa Abd El Meguid
Centre for Intelligent Machines, McGill University
Montreal, Canada
Date: 26/04/2010

Symmetries of Occluded and Non-Rigid Objects

Mostafa Abd El Meguid
Centre for Intelligent Machines, McGill University
Montreal, Canada

I. INTRODUCTION

Symmetry is an ever present geometric phenomenon. Nature is abundant with examples of symmetric structures down to the microscopic level. Inspired by what was around them, humans went on to extensively include symmetry into their manufactured products.

One reason for the prevalence of symmetry in both natural and artificial domains is the simple aesthetic appeal it carries. Gestalt laws of proximity, similarity, common fate, and objective set [1] can all attribute some of their significance to inherently present symmetric properties. The adherence to said gestalt laws, or ‘laws of perceptual form’ in human discernment speaks to the biological importance of symmetry in representations of objects in the brain.

Academic studies of symmetry date back to Ancient Greece [2], mostly discussing to the beauty of measured proportion in sculptures. The building of a mathematical basis for symmetry did not quite begin till Johannes Kepler discussed crystals in his 1611 study “*Strena seu de nive sexangula*” or “*On the Six-Cornered Snowflake*” [3].

Generally labeling symmetry as ‘self-similarity’ brings about a sense of vagueness. However, as Hermann Weyl alluded to in [2], symmetry can be as widely or narrowly defined. An expansion on Weyl’s statements is that symmetry can be defined on various scales, collectively or individually.

Different scales relates to how much an object is broken down before the various parts are analyzed for symmetry. Grouping relates to how the aforementioned parts contribute to their parent’s symmetric measure. The problem of identifying a standalone object or what defines its parts is not explicitly handled in the algorithms discussed and implemented here, although, the significance of part symmetry will become quite obvious.

This project involves a study of two seemingly disjoint algorithms. The first algorithm discusses different types of three dimensional geometric symmetries, which are then used to complete partially occluded models in a stochastic manner. The second algorithm proposes a technique to measure the symmetric properties of non-rigid shapes, utilizing the geodesic metric on a mesh representing a three dimensional object.

II. SHAPE FROM SYMMETRY [5]

Thrun and Wegbreit propose a three dimensional object model completion algorithm. The basis of the algorithm is customized taxonomy of symmetry features. Through the hierarchical nature of said taxonomy, an entailment of symmetries arises.

Reconstruction of partially occluded segments in images using symmetry is a topic that been exhaustively studied as summarized in [6]. Systems employing the continuous symmetry measure [7] have shown some success in reconstructing complex partially symmetric objects as shown in Figure 3.

The inspiration for this paper comes from the common presence of symmetry in natural and artificial domains as mentioned earlier. Analogous symmetry features exist between \mathbb{R}^2 and \mathbb{R}^3 , providing a stepping stone for expansion of the two dimensional reconstruction algorithm.

Three dimensional models obtained using active stereo rigs may be partially occluded. The occlusion may arise due to the relative placement of the scanner and the object, thus the object’s ‘front’ occludes its ‘back’. Should the object possess symmetric features, the proposed algorithm would be able to detect and utilize their geometric arrangement, along with the scanner’s location, to extend the viewable segment into the occluded space.

It is worthwhile to note that the scanner captures various surface characteristics (convexities, concavities, hyperbolic and parabolic regions) as opposed to imaging cameras that fail to capture concavities of surfaces.

A. Types of Symmetry

Three dimensional symmetry features are defined using three Euclidean geometry primitives: planes, lines and points, all lying in \mathbb{R}^3 . These three primitives act as ‘reflection symmetries’ that associate a pair of points, P and Q , on the surface of an object as shown in Figure 4.

Manifolds, Q , are associated with a surface point P . In the case of the manifold (note that $P \in Q$) forming the circumference of a circle lying in a plane, the center and normal of said circle and plane, respectively, define an axis of reflection. Figure 6 (a) shows an example of axial symmetry. In the case of the manifold Q forming a spherical surface, the center of the sphere defines a spherical symmetry point feature as shown in Figure 6 (b).

Combining multiple symmetry features, such as two or three planes, extends the pair association to a point set correspondence. Figure 7 (b) shows the case of two planes where ‘Plane A’ associates points (P, Q_1) and (Q_2, Q_3) , while ‘Plane B’ associates points (P, Q_2) and (Q_1, Q_3) . Together, ‘Plane A’ and ‘Plane B’ build the association (P, Q_3) and (Q_1, Q_2) , resulting in P being associated with the point set $Q = \{Q_1, Q_2, Q_3\}$. When three planes are combined, as shown in Figure 7 (c), the point set Q consists of 7 individual points, while in the case of combination of a plane and an axis, Q consists of two manifolds forming independent circles with aligned centres.

This list of symmetries is limited to what the authors deemed the most prevalent in natural environments.

B. Entailment Hierarchy

Figure 1 shows a hierarchy of the symmetries defined above. The hierarchy is meant to be transcended top-down, such that a spherical symmetry automatically implies the existence of combined axial, planar reflection symmetry or a combination of three reflection planes etc.

This hierarchy reduces the search space at any stage of the algorithm as shown later on. Intuitively, if no planar symmetry is found, no dual planar symmetry can exist, i.e. when transcending the hierarchy bottom-up, each symmetry (box) acts as a gateway to the next. Since some symmetry feature combinations neither entail nor are entailed by other symmetries, they cannot be included in the hierarchy.

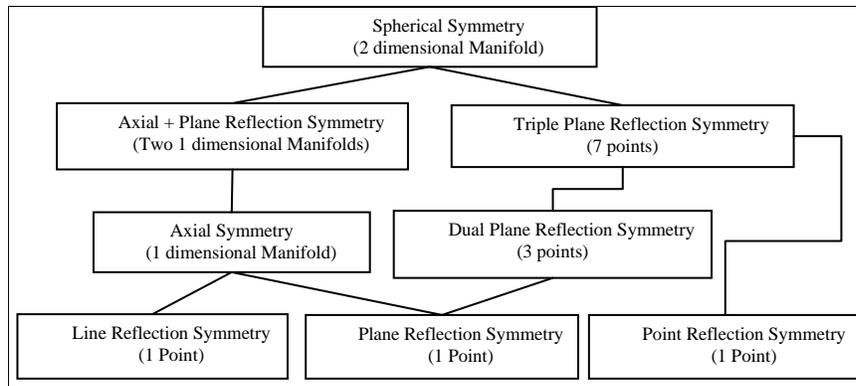


Figure 1. The Entailment Hierarchy of Symmetries [5]

C. Probabilistic Measurement Model

The development of a probabilistic model is strongly dependant on the definition of occluded and unoccluded space with respect to the scanner location.

Consider a potential reflected point set Q coming from an original point set \mathcal{P} generated by a 3D scanner. With the occluded and unoccluded views of the scanner well-defined (whether through aperture width or object presence) the point set Q is segmented into a portion that lies in said occluded view Q_{occl} , and a portion that lies in the unoccluded view, Q_{match} . Thus the main restriction permitting Q to be a valid set is that Q_{match} should lie entirely in \mathcal{P} otherwise the reflection becomes implausible.

Naturally, any imaging process, such as 3D scanning, will result in a noisy, discretised point set. To accommodate these limitations, a probabilistic model is developed using two Gaussian distributions. First, the

probability that a point $Q \in \mathcal{Q}$ came from the visible point set, \mathcal{P} , due to the corresponding coordinates of Q and a point $P \in \mathcal{P}$, is defined as shown in (1).

$$p(Q|\mathcal{P}, \text{match}) = \arg \max_{P \in \mathcal{P}} \frac{1}{\sqrt{2\pi|\Sigma_{\text{match}}|}} e^{-\frac{(Q-P)^T \Sigma_{\text{match}}^{-1} (Q-P)}{2}} \quad (1)$$

Equation (2) models the probability that a point $Q \in \mathcal{Q}$ is a reflection of a point, \bar{P} , in the point set, $\bar{\mathcal{P}}$ representing the occluded portion of the object (obtained using ray casting)

$$p(Q|\mathcal{P}, \text{occl}) = \arg \max_{\bar{P} \in \bar{\mathcal{P}}} \frac{1}{\sqrt{2\pi|\Sigma_{\text{occl}}|}} e^{-\frac{(Q-\bar{P})^T \Sigma_{\text{occl}}^{-1} (Q-\bar{P})}{2}} \quad (2)$$

To model anomalies in \mathcal{Q} , a constant η is used as shown in (3), where η is equal to the inverse of the maximum measurement range.

$$p(Q|\mathcal{P}, \text{random}) = \eta \quad (3)$$

Using Bayes' rule, the conditional probabilities in (1), (2) and (3) are combined to form the probability that a potential reflected point, Q , was obtained from the visible point set \mathcal{P} .

$$p(Q|\mathcal{P}) = p(\text{match}) p(Q|\mathcal{P}, \text{match}) + p(\text{occl}) p(Q|\mathcal{P}, \text{occl}) + p(\text{random}) p(Q|\mathcal{P}, \text{random}) \quad (4)$$

Note that $p(\text{match}) + p(\text{occl}) + p(\text{random}) = 1$. An approximation to the logarithm of $p(Q|\mathcal{P})$ is shown below in (5).

$$-\log p(Q|\mathcal{P}) = \text{constant} + \frac{1}{2} \left[p(\text{match}) \arg \max_{P \in \mathcal{P}} (Q-P)^T \Sigma_{\text{match}}^{-1} (Q-P) + p(\text{occl}) \arg \max_{\bar{P} \in \bar{\mathcal{P}}} (Q-\bar{P})^T \Sigma_{\text{occl}}^{-1} (Q-\bar{P}) + p(\text{random}) \log \eta \right] \quad (5)$$

Observe that the first bracketed term in (5) corresponds to the quadratic surface distance from the iterative closest point algorithm. The second term may not intuitively belong in the expression; however, the point of the algorithm is to reconstruct an occluded object by extending the 'visible' point set into the occluded space.

The expression shown in (5) is extended within the potential point set, \mathcal{Q} , assuming independence from point to point. This extension provides a quantitative scoring mechanism for a proposed point set as shown in (6).

$$-\log p(Q|\mathcal{P}) = -\frac{1}{|\mathcal{Q}|} \sum_{Q \in \mathcal{Q}} \log p(Q|\mathcal{P}) \quad (6)$$

D. Searching for Symmetries

The entailment hierarchy shown in Figure 1 is used to constraint the search space at each point. Not only do lower level symmetries permit or restrict higher level symmetry features; they also constrain the parameters defining said higher level symmetry features.

The search for symmetries utilizes the constraints given by the entailment hierarchy and enumerates plausible symmetry parameters defining various symmetry features at each point. Once a symmetry feature is deemed plausible, the set of points, \mathcal{Q} , arising from said features are calculated and their collective score is obtained using the expression in (6). A gradient based optimization method is then used to tune the symmetry features possessing minimal score.

A heuristic evaluation is then used to segment surface portions into 'parts'. In this setting, each part strongly suggests the existence of a particular symmetry feature that may not be globally prevalent.

E. Observations

The results shown in the paper are quite positive, displaying successful segmentation of various parts depending on the symmetry features suggested by the constituent points. However, several questions require answering before the algorithm can even be implemented.

1) Prior Probabilities

The probabilities $p(\text{match})$, $p(\text{occl})$, and $p(\text{random})$ used in (5) are specific to the 3D point cloud and scanner position considered. The values attached to each of the prior probabilities would greatly affect the scoring mechanism employed in this paper. It is expected that $p(\text{random})$ would be quite small relative to the

other two priors, however, a near exhaustive search would be required to determine a globally appropriate value for $p(\text{match})$ and $p(\text{occl})$. One heuristic that could be used in this step is the dimension of the initial visible set.

2) *Part Segmentation*

The heuristic employed to segment the object according to the suggested symmetries uses the implied volume enclosed by the surface forming said segment. The relation between the surface and the volume is straightforward for basic shapes with successfully discovered surfaces. However, once surface patches are incorrectly suggested, the volume-surface area relation becomes invalid. Thus, for successful part segmentation, a correct symmetry feature extraction stage is required.

3) *Camera Location*

Should the two above limitations be successfully circumvented, the algorithm still requires a specific camera location to be paired with the provided point set. This requirement could be removed; however, the search space would have to be expanded significantly, rendering the algorithm quite expensive.

III. SYMMETRIES OF NON-RIGID SHAPES [4]

The algorithm described in [5] and in Section II provides a measure of symmetry based on the relationships between points in \mathbb{R}^3 . This technique is useful to discover the existence of symmetry in non-deformable or rigid objects (i.e. objects whose parts cannot move independently, thus Euclidean distances between different points on the object surface remain the same). However, Euclidean geometric measures are not as robust in detecting regional symmetries in objects whose parts can move freely, breaking global symmetries in the process. To detect such symmetries with better accuracy, an algorithm relying on surface bound (geodesic) distances is developed in [4] and described below.

A. *Extrinsic and Intrinsic Symmetries*

1) *Extrinsic Symmetries*

Consider the example of the human model shown in Figure 8 (a), a single global symmetry feature exists (reflection plane) that pairs every two surface points in a bijective nature. This is a type of extrinsic symmetry defined by the extrinsic geometry of the object.

A shape, X , consisting of a three dimensional point set can be considered a metric space under the restricted Euclidean metric $d_{\mathbb{R}^3|X}: X \times X \rightarrow \mathbb{R}$. The restricted Euclidean metric measures the shortest path between every pair of surface points of X in \mathbb{R}^3 , in the process defining the extrinsic geometry of X .

Defining a transformation, $g: X \rightarrow \mathbb{R}^3$ mapping X onto itself, such that $g(X) = X$, while preserving the extrinsic geometry of X , as a self-congruence or an extrinsic symmetry, gives the following function composition for the group of extrinsic symmetries X :

$$\text{Esym}(X) = \{g: X \xrightarrow{1:1} X : d_{\mathbb{R}^3|X} = d_{\mathbb{R}^3|X} \circ (g(X) \times g(X))\} \quad (7)$$

Observe that $\text{Esym}(X)$ contains the notions of reflection and rotation symmetries mentioned in Section II.A.

2) *Intrinsic Symmetries*

Consider the example of the human model shown in Figure 8 (b) where the single global plane of symmetry in Figure 8 (a) no longer applies, however, part symmetries still exist (head, legs, arms etc..). These part symmetries also relate to the reflections and rotations defined in Section II.A, however, as mentioned in Section II.E, discerning the individual parts and the symmetries they support is not a straightforward process.

Equipping X with the geodesic metric $d_X: X \times X \rightarrow \mathbb{R}$ measuring the shortest path between every path of surface points of X on the mesh representing X allows for the definition of the intrinsic geometry of X .

Defining a transformation $g: X \rightarrow \mathbb{R}^3$ mapping X onto itself, such that $g(X) = X$, while preserving the intrinsic geometry of X , as an isometry or an intrinsic symmetry, gives the following function composition for the group of intrinsic symmetries of X :

$$\text{Isym}(X) = \{g: X \rightarrow X : d_X = d_X \circ (g(X) \times g(X))\} \quad (8)$$

B. Mappings and Symmetries

Let $\mathbb{F}(X)$ as the space of mappings (or transformations) $g: X \rightarrow X$. Without loss of generality, define the metric $d_{\mathbb{F}(X)}(f, g)$ on $\mathbb{F}(X)$, for all $f, g \in \mathbb{F}(X)$ as shown in (9).

$$d_{\mathbb{F}(X)}(f, g) = \max_{x \in X} d_X(f(x), g(x)) \quad (9)$$

Defining a closed metric ball of radius r centered at g as $B_{\mathbb{F}(X)}(g, r)$ as shown in (10).

$$B_{\mathbb{F}(X)}(g, r) = \{f \in \mathbb{F}(X): d_{\mathbb{F}(X)}(g, f) \leq r\} \quad (10)$$

Furthermore, associate the distortion $\text{dis}(g)$ with every mapping $g \in \mathbb{F}(X)$ as shown in (11). $\text{dis}(g)$ measures the difference between d_X and $d_X \circ (g \times g)$.

$$\text{dis}(g) = \max_{x, x' \in X} |d_X(x, x') - d_X(g(x), g(x'))| \quad (11)$$

Combining the definition of $\text{ISym}(X)$ from (8) with that of the distortion in (11) gives the ideal condition for an intrinsic symmetry to be $\text{dis}(g) = 0$. Thus, intrinsic symmetries can be seen as global minima of the distortion space. Furthermore, intrinsic symmetries can be found to be local minima of the distortion such that for any $f \in B_{\mathbb{F}(X)}(g, r)$, where r is a sufficiently small radius, $\text{dis}(f) \geq \text{dis}(g)$.

When dealing with a discretised (through acquisition) point set X , chances that a non-identity (non-trivial) transformation possesses zero distortion are quite low. Thus, the approximate symmetry set (ϵ -symmetry set), $\text{ISym}(X, \epsilon)$, can be defined as shown in (12), where and appropriate ϵ is to found through experimentation.

$$\text{ISym}(X, \epsilon) = \left\{ g \in \mathbb{F}(X): \epsilon \geq \text{dis}(g) = \min_{f \in B_{\mathbb{F}(X)}(g)} \text{dis}(f) \right\} \quad (12)$$

C. Evaluating Symmetries

In contrast to the scoring scheme employed in Section II.C, this algorithm uses an asymmetry measure to score a mapping $g \in \mathbb{F}(X)$. The local shape asymmetry is defined at every point $x \in X$ as shown in (13). The local shape asymmetry measures the distortion of g at x .

$$\text{asym}(X, x) = \max_{x' \in X} |d_X(x, x') - d_X(g(x), g(x'))| \quad (13)$$

Maximizing across all local asymmetries gives the global asymmetry measure shown in (14).

$$\text{asym}(X) = \max_{x \in X} \text{asym}(X, x) \quad (14)$$

D. The Algorithm

Given a large point set, X , representing an object, the ‘farthest point sampling algorithm’ is used to obtain an r -covering sampled set, X_r , consisting of N points. The farthest point sampling algorithm guarantees the geodesic separation between any two points is greater than the sampling radius r , i.e. $d_X(x_i, x_j) \geq r$ with $x_i, x_j \in X_r, i \neq j$. X_r is approximated using a mesh, \hat{X} , that constrains the space of mappings $\mathbb{F}(\hat{X})$. The asymmetry measures shown in (13) and (14) are calculated across $\mathbb{F}(\hat{X})$ in an attempt to extract the mapping with minimal distortion, corresponding to the most probable isometric transformation.

The algorithm proceeds in two broad steps. First, a coarse initial estimate of the minimum distortion mapping is obtained using a sub-sampling of the set X_r . The estimate is then fine tuned using a multidimensional scaling scheme.

1) Initialization

X_r is first sub-sampled with a radius $R \gg r$, resulting in point set X_R containing M points such that $M \ll N$. The set of mappings $\mathbb{F}(X_R)$ corresponds to all discrete permutations of X_R such that $\pi: X_R \rightarrow X_R$. Removing the identity mapping from $\mathbb{F}(X_R)$ leaves an $\mathcal{O}(M^M)$ mapping distortion calculation to obtain a minimal mapping. This search space is reduced through the discarding of mappings where the metric d_X behaves differently at the original point x_i and its transformed point π_i .

The process of discarding mappings begins with the construction of a histogram of geodesic distances across the set X_R as shown in (15), where $\hat{d}_{ij} = d_{X_R}(x_i, x_j)$.

$$h_i = \text{hist}(\{\hat{d}_{ij}: \hat{d}_{ij} \leq \rho\}) \quad (15)$$

The threshold ρ in (15) corresponds to the radius of the sphere about $x_i \in X_R$ in which the geodesic distances are calculated. The set of histogram vectors are used as descriptors for their relative points providing a discrepancy measure, $d(h_i, h_j)$, between points as shown in (16). The discrepancy measure is a weighted histogram distance where the distance between bins in $h_i - h_j$ is accounted for using the matrix A_{mn} .

$$d(h_i, h_j) = \sqrt{(h_i - h_j)^T A (h_i - h_j)} \quad (16)$$

After calculation of $d(h_i, h_j)$ for all points in X_R , a set of the K ‘most likely’ transform points is constructed for each point $x_i \in X_R$. Labeling the set of most likely transform points C_i , the space of mappings $\mathbb{F}(X_R)$ can be reduced to a new space $\mathbb{F}_{\text{init}} = C_1 \times C_2 \times \dots \times C_M$.

Exhaustively searching the space \mathbb{F}_{init} remains an expensive process, thus, a ‘hierarchical greedy algorithm’ is employed to extract a ‘good enough’ preliminary mapping.

a) *Pairing:*

Taking a pair (i, j) from the set X_R , find the pair $(m, n) \in C_i \times C_j$ minimizing the distortion $|\hat{d}_{ij} - \hat{d}_{mn}|$. This process gives a set of two point correspondences, E_2 , sorted in increasing order of distortion.

b) *Merging:*

Each two pairs in E_2 are in turn merged into a four pair set based on another distortion minimization scheme. The one restriction is the requirement of all points in the set being unique. This process forms a new set E_4 and stops when E_2 is empty. Repeatedly merging to produce E_{2k} from E_k , and halting once E_{16} is obtained.

c) *Completion:*

Given a set of k point correspondences, E_k , such that $(i_1, \dots, i_k) \mapsto (\pi_1, \dots, \pi_k)$ is a single k point correspondence. Each k point correspondence is then extended to the remaining points in X_R as shown in (17).

$$\pi_j = \arg \min_{\pi_j \in \{x_1, \dots, x_M\}} \max_{i \in \{i_1, \dots, i_k\}} |\hat{d}_{ij} - \hat{d}_{\pi_i, \pi_j}|, j = k + 1, \dots, m \quad (17)$$

The set of candidate transformations are then analyzed for distortion and the minimal mapping π_{min} is obtained. The unidirectional nature of the algorithm does not guarantee that a minimal mapping is found, however, experiments show that the result is highly dependent on the pairs forming E_2 . Note that the initial pairs are selected in spite of potential draws in distortion minimization, thus the order of selection of pairs would have an effect on the eventual mapping obtained.

2) *Local Refinement*

A multiresolution optimization scheme is employed to obtain the mapping with minimum distortion from the coarse estimate, π_{min} . Using the notation $x'_i = g(x_i)$, the minimizing set of (18) represents the minimal mapping.

$$\min_{\{x'_1, \dots, x'_N\} \subset \mathbb{F}(\hat{X})} \max_{i, j=1, \dots, N} |\hat{d}_{ij} - \hat{d}_X(x'_i, x'_j)| \quad (18)$$

Observe that the refined minimizing mapping corresponds to minimizing over the entire mesh \hat{X} . The minimizing set $\{x_1^*, \dots, x_N^*\}$, found using convex optimization techniques is used to obtain approximations for (13) and (14) on the mesh \hat{X} as shown in (19) and (20).

$$\text{asym}(\hat{X}, x_i) = \max_{j=1, \dots, N} |\hat{d}_{ij} - \hat{d}_X(x_j^*, x_i)| \quad (19)$$

$$\text{asym}(\hat{X}) = \max_{i,j=1,\dots,N} |\hat{d}_{ij} - \hat{d}_X(x_i^*, x_j^*)| \quad (20)$$

E. Implementation

1) Preliminary Notes:

a) Toolbox References

The implementation here makes use of the MATLAB toolboxes in [9] and [12] which implement several algorithms related to geodesic mesh traversal and distortion calculations.

b) Data Sources

The 3D mesh models used in this project came from two data sets, [14] and [15].

2) Goal

The algorithm implemented here will attempt to discover the intrinsic symmetry of certain pose of a non-rigid object relative to an intrinsically symmetric version of the same shape, which acts as a reference. Although the use of a reference may seem trivial at first, it will become clear as the results emerge and the potential uses are discussed.

The expected outcome is a relative score of symmetry within a set of poses with a fixed reference. The different steps of the algorithm up to the Local Refinement stage (Section III.D.2) were implemented in MATLAB while the multiresolution optimization process was carried out using the scheme described in [11] and implemented in [9].

The code shown in **Error! Reference source not found.** implements the steps in Section III.D.

F. Results

The refinement (or mapping extension) was carried out with $M = 50$ and $N = 200$ rather than the full mesh simply due to computational and temporal limitations. The reduced workspace will simply utilise a smaller set of correspondences which would still be representative of the overall point set but with a larger r -covering radius. Voronoi tessellation was then used to group points according to their proximity to one of the selected refinement points.

The reference pose is shown in Figure 2. The first column of relates to the reference pose used to base symmetry scoring. The columns of Table 1 relate to varying poses with increasing qualitative asymmetry (decided via personal opinion). The first row shows a smooth mesh representation of the pose. The second row shows the color coded Voronoi tessellation of the ‘front’ of the posed object and the third row shows the corresponding tessels on the original pose. The fourth row shows an intensity image based on the local shape asymmetry in (19), (red corresponds to high level of asymmetry). The fifth row shows the global asymmetry value of the pose, calculated as shown in (20).



Figure 2. Reference Pose used in Analysis shown in Table 1

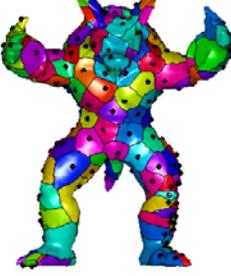
Pose:	Pose 1	Pose 2	Pose 3	Pose 4	Pose 5	Pose 6
Smooth Mesh Representation						
Pose Voronoi Tessellation						
Corresponding Reference Voronoi Tessellation						
Intensity Image of Local Asymmetry						
Global Asymmetry	22.1468	24.3028	26.2115	28.3975	29.2869	57.9710

Table 1 Results of running the Non-Rigid Symmetry Detection Algorithm

G. Discussion

The first observation regarding the results in Table 1 is that the algorithm mostly succeeds at segmenting surface patches and correctly relating them to their corresponding reference. Limitation arise at patches corresponding to surface regions that have changed dimensions. This is due to the mesh possessing an ‘elastic skin’. The change in dimension corresponds to a change in d_X , the metric off of which the entire algorithm is based. However, the algorithm is still able to accommodate relatively major changes in d_X as shown by the knee bending of Pose 3.

Pose 4 is an interesting case in terms of the tessellation correspondence. Observe that the colors of the patches on the left leg of the deformed object correspond to the right leg of the reference object, similarly for the right leg of the deformed object and the left leg of the reference. The reason for this switch is the similar geodesic distance map between the two legs of the object. Tracing this correspondence of a point on the left deformed leg to the right reference leg through the algorithm, it was seen that the pairing and merging algorithm was convinced that a point on the right leg of the reference object experienced a similar d_X mapping. A point on the left reference leg was also selected in the C group, however, the random selection process ended up picking the opposite leg.

The intensity plot visualising the local shape asymmetry is quite representative of the difference between the pose considered and the reference object. Regardless of orientation or location in \mathbb{R}^3 , if no change has occurred in d_X over a surface region in the pose relative to the reference, a low local asymmetry value is obtained. This is best exemplified by the torso region in Pose 3. Even with arms flailing and legs knelt, the upright torso was quickly matched to the reference torso with a small level of asymmetry.

Finally, considering Pose 6, where the Armadillo is seated, altering a large portion of the surface’s d_X map. The algorithm was unable to find a good set of initial point correspondences with the sampling quantities employed. A larger sampling might be able to find a better set of initial mappings to use in the GMDS local refinement step.

H. Conclusion

Overall, the algorithm succeeded in discovering regions with a high level of discrepancy in d_X relative to a reference object. The failure came in an extreme case (Pose 6 in Table 1) which does not represent the ideal application for the algorithm.

The ability of this algorithm to accommodate non-rigid shapes, along with its simple 3D visualisation and quantization of asymmetry makes it ideal for quality control stages in industrial processes. It is here that the relative score comes into light. Given an ideal model of a product, an acceptable asymmetry score threshold is set, judging the validity of a prospective product according to its representative 3D model. Similarly, with 3D medical modelling developing into an area of increased research, medical diagnosis expert systems could benefit from the algorithm in identifying unhealthy tissues given a healthy model.

IV. REFERENCES

- [1] Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt II. *Psychologische Forschung*, 4, 301-350.
- [2] Weyl, H. (1983). *Symmetry*. Princeton University Press
- [3] Kepler, J. (1611). *Strena seu de nive sexangula*. Frankfurt.
- [4] Raviv, D., Bronstein, M. M., Bronstein, A. M., & Kimmel, R. (2007). Symmetries of non-rigid shapes. *Eleventh IEEE International Conference on Computer Vision (ICCV'07)*, (pp. 1-7). Rio de Janeiro.
- [5] Thrun, S., & Wegbreit, B. (2005). Shape From Symmetry. *Tenth IEEE International Conference on Computer Vision (ICCV'05)*, 2, pp. 1824-1831. Beijing.
- [6] Zabrodsky, H., Peleg, S., & Avnir, D. (1993). Completion of Occluded Shapes Using Symmetry. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR '93)*, (pp. 678-679). New York City
- [7] Zabrodsky, H., Peleg, S., & Avnir, D. (1995, December). Symmetry as a Continuous Feature. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17, pp. 1154-1166.
- [8] Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2007, September/October). Calculus of non-rigid surfaces for geometry and texture manipulation. *13 (5)*, pp. 902-913.
- [9] Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2009, 1 10). *CS468 Topics in Geometric Algorithms (Numerical geometry of non-rigid shapes)*. Retrieved 4 20, 2010, from Project TOSCA: http://tosca.cs.technion.ac.il/book/course_stanford09.html

- [10] Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006). Efficient computation of isometry-invariant distances between surfaces. *SIAM Journal Of Scientific Computing* , 28/5, 1812-1836.
- [11] Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006). Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of National Academy of Sciences (PNAS)*, 103/5, pp. 1168-1172.
- [12] Peyre, G. (2008). *Toolbox Fast Marching - A toolbox for Fast Marching and level sets computations*. Retrieved 04 20, 2010, from <http://www.ceremade.dauphine.fr/>: <http://www.ceremade.dauphine.fr/~peyre/matlab/fast-marching/content.html#33>
- [13] Sethian, J. A. (1999). *Level Set Methods and Fast Marching Methods*. Cambridge University Press.
- [14] Stanford Computer Graphics . (2006, 04 05). *Armadillo Deformations*. Retrieved 04 20, 2010, from AIM at Shape - Digital Shape Workbench: <http://shapes.aim-at-shape.net/viewgroup.php?id=657>
- [15] TOSCA. (2009, 01 10). *TOSCA high-resolution Data Set*. Retrieved 04 20, 2010, from Project TOSCA: http://tosca.cs.technion.ac.il/book/resources_data.html

V. APPENDIX I: EXTERNAL FIGURES

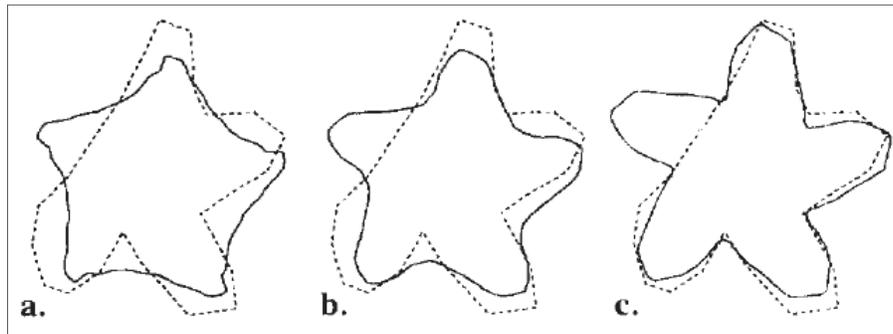


Figure 3. Reconstruction of an occluded almost symmetric shape. The original shape (dashed line) and the reconstructed shape (solid line). The closest symmetric shape following angular selection about: a) the centroid, b) the center of symmetry, c) the center of symmetry using clustering in the completion [6]

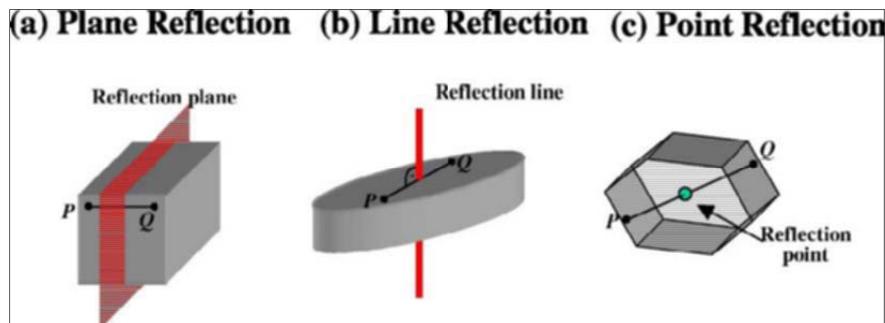


Figure 4.

Figure 5. Reflection Symmetries [5]

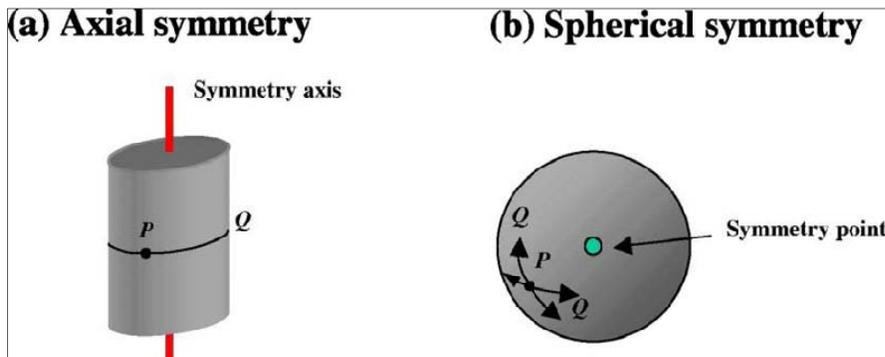


Figure 6. Axial and Spherical Symmetries [5]

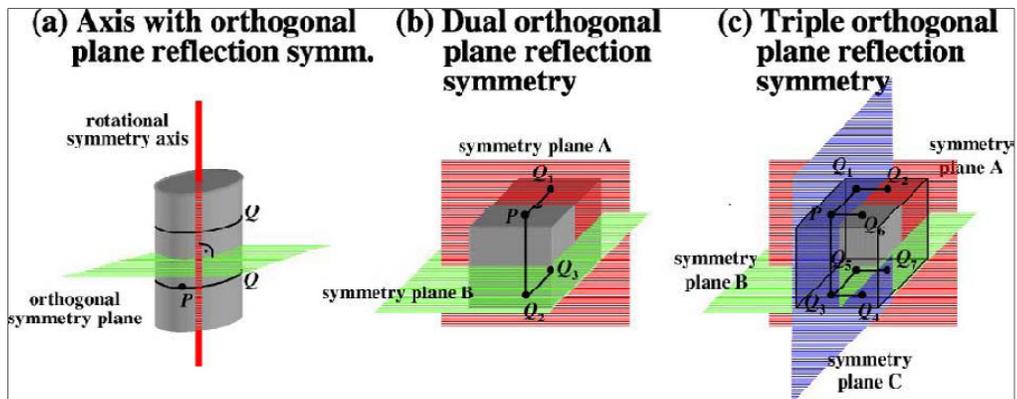


Figure 7. Composite Symmetries[5]

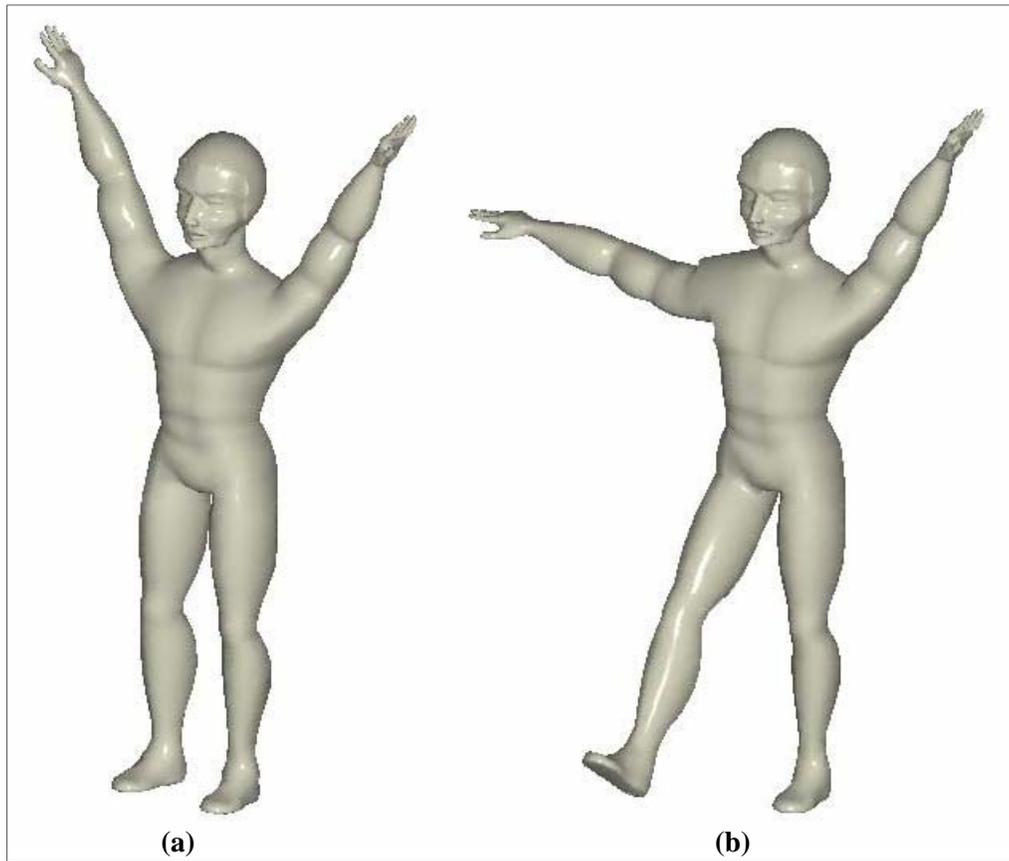


Figure 8. Visualization of the difference between Extrinsic and Intrinsic Symmetry. (a) Extrinsically Symmetric shape is also intrinsically symmetric, (b) Isometry of the shape is Intrinsically Symmetric but Extrinsically Asymmetric [4]