THE HEBREW UNIVERSITY OF JERUSALEM
INSTITUTE OF COMPUTER SCIENCE

TR #94-2 December, 1993

Three Dimensional Symmetry from Two Dimensional Data

Hagit Zabrodsky and Daphna Weinshall

Abstract

In previous applications, bilateral symmetry of objects was used either as a descriptive feature in domains such as recognition and grasping, or as a way to reduce the complexity of structure from motion. In this paper we propose a novel application, using the symmetry property to "symmetrize" data before and after reconstruction. We first show how to compute the closest symmetric $2D$ and $3D$ configurations given noisy data. This gives us a symmetrization procedure, which we apply to images before reconstruction, and which we apply to the $3D$ configuration after reconstruction. We demonstrate a significant improvement obtained with real images. We demonstrate the relative merits of symmetrization before and after reconstruction using simulated and real data.

This report describes research done at the Hebrew Univresity of Jerusalem within the Institute of Computer Science. This research was sponsored by the U.S. Office of Naval Research under Grant N00014-93-1-1202, R&T Project Code 4424341—01. DW was also supported by the Israeli Science Foundation grant 202/92-2.

1 Introduction

The most common symmetry in our environment is three dimensional mirror symmetry. It is thus not surprising that the human visual system is most sensitive to bilateral symmetry. A common case in human and computer vision is that only $2D$ (projective) data is given about a $3D$ object. Many studies deal with inferring $3D$ symmetry from $2D$ data. These studies deal with perfect non-noisy data. In this paper, we deal with noisy $2D$ data by extending the notion of Symmetry Distance defined in [17] to $2D$ projections of $3D$ objects which are not necessarily perfectly symmetric. We describe in this work the reconstruction of $3D$ mirror symmetric connected configurations from their noisy $2D$ projections.

More specifically, we describe the enhancement in performance that can be obtained using existing structure from motion (or structure from a sequence of $2D$ images) methods, when the reconstructed object is known to be mirror-symmetric. We consider here objects whose $3D$ structure is a mirror-symmetric connected configuration (a $3D$ graph structure composed of one or more connected components). We are given several noisy $2D$ projections of such an object, where the projection is approximately weak perspective (scaled orthographic). In this work we combine the invariant reconstruction algorithm described in [16] with a symmetrization method, to improve the input and output data in the structure reconstruction from several views.

We employ two approaches to exploit the fact that the $3D$ structure to be reconstructed is mirror-symmetric:

- correct for symmetry prior to reconstruction
- correct for symmetry following reconstruction

Correction for symmetry following reconstruction is performed by applying any existing method of structure from motion with no a-priori symmetry assumption on the reconstructed object. Following the reconstruction, the symmetry assumption is exploited and the mirror-symmetric structure closest to the reconstruction is found. This last stage is performed using a closed form method described in Section 3 for finding the closest mirror-symmetric configuration to a given $3D$ connected configuration.

Correction for symmetry prior to reconstruction requires application of some symmetrizing procedure to the $2D$ data with respect to the $3D$ symmetry. Following the symmetrization procedure, any existing method of reconstruction of general $3D$ structure from $2D$ data can be applied. Notice that this procedure does not ensure that the final reconstructed $3D$ structure is mirror-symmetric; however, as will be shown in Section 5, the error in reconstruction is greatly reduced. In Section 4 we describe a symmetrization procedure of $2D$ data for projected $3D$ mirror-symmetry. In Section 5 we give examples and comparisons between correction for symmetry prior and following $3D$ reconstruction, using real and simulated data.

2 Previous work

As an intrinsic characteristic of objects and shapes, symmetry can be used to describe and recognize objects. Many studies deal with symmetry of $2D$ shapes and patterns, either using global symmetry features [6, 2] or local symmetry features [4, 11]. Most studies deal with a specific symmetry class such as circular (radial) symmetry [2, 11] or mirror and linear symmetry [8, 5]. A method for dealing with any type of symmetry in both $2D$ and $3D$ has recently been suggested [19, 18, 17].

When dealing with $3D$ symmetries, several studies concentrate on finding the projected or skewed symmetries in $2D$ images [5, 10]. Other studies reconstruct $3D$ objects from $2D$ images using symmetry as a constraint [13, 14]. Symmetrical descriptions and symmetric features of objects are useful in guiding shape matching, model-based object matching and object recognition [8, 12]. Additionally, symmetry has been used in guiding robot grasping [3]. Recently, symmetry has been exploited for reducing complexity and reducing the number of frames in structure from motion problems [9, 7, 12]. However none of these studies deal with exploiting symmetry for improving the input data for structure from motion algorithms, or the symmetrization of noisy two dimensional projections of three dimensional obejcts, which is the subject of the present paper.

3 3D symmetrization

In [18, 17] we described a method for finding the symmetric configuration of points which is closest to a given configuration in a least squares sense. We defined a measure of symmetry - the Symmetry Distance (SD), and described a method for evaluating this measure for any configuration of points with respect to any point symmetry group in any dimension. An outcome of evaluating the Symmetry Distance of a given configuration is the configuration which is symmetric and which is closest to the original configuration in a least squares sense. An *iterative* folding/unfolding method, which finds the closest symmetric configuration, was described in [18, 17]. Below we describe another *closed-form* solution that gives equivalent results in the case of $3D$ mirror-symmetry.

We first note that every mirror symmetric $3D$ configuration of points $\{P_i\}_{i=0}^{n-1}$ implicitly implies a pairing (matching) of the points: for every point P_i there exists a point $match(P_i) = P_j$ which is its counterpart under reflection. Following is the closed-form algorithm as applied to $3D$ mirror symmetry (Fig. 1).
Given a configuration of points $\{P_i\}_{i=0}^{n-1}$ in $\mathcal{R}^3$ (see Fig. 1a):

1. **Divide** the points into sets of one or two points. If a set contains one point, duplicate that point. In the example of Fig. 1, the sets are $\{P_0, P_0\}, \{P_1, P_3\}$ and $\{P_2, P_2\}$. This defines a matching on points of the object.

2. **Reflect** all points across a randomly chosen mirror plane, obtaining the points $\tilde{P}_i$ (Fig. 1b).

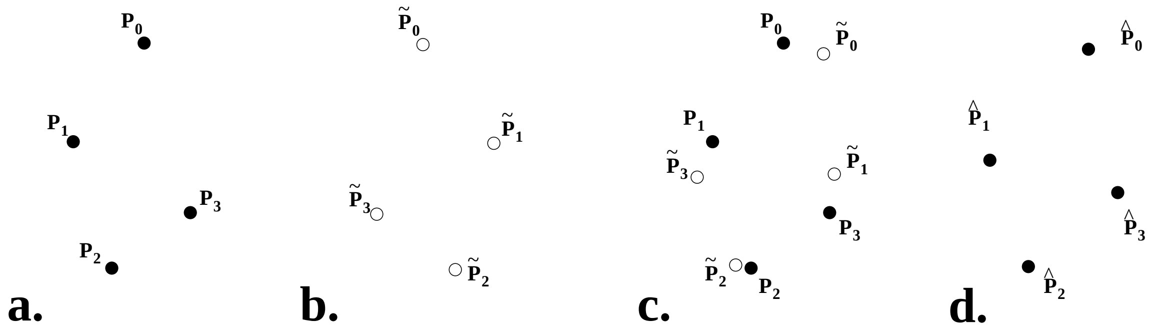

Figure 1: Obtaining the closest mirror symmetric set of points - see text.

3. **Find** the optimal rotation and translation which minimizes the sum of squared distances between the original points and the reflected corresponding points (Fig. 1c). This is a well known problem of pose estimation. To find the solution we use the method of Arun et. al. [1], which requires no more than the evaluation of SVD.

4. **Rotate** the reflected points $\{\tilde{P}_i\}_{i=0}^{n-1}$ by the optimal rotation.

5. **Average** each original point P_i with its reflected matched point $\tilde{P}_j$, obtaining the point $\hat{P}_i$ (Fig. 1d). The points $\{\hat{P}_i\}_{i=0}^{n-1}$ are mirror symmetric.

6. **Evaluate** the symmetry distance value: $\frac{1}{n}\sum_{i=0}^{n-1} \|P_i - \hat{P}_i\|^2$.

7. **Minimize** the symmetry distance value obtained in step 6 by repeating steps 1-6 with all possible division of points into sets. The mirror symmetric configuration corresponding to the minimal symmetry value is the closest mirror symmetric configuration in a least squares sense (proof is given in [17]).

In practice, the minimization in step 7 is greatly simplified when the configuration of points is connected (or partially connected). Consider the original configuration as a graph $G = \{V, E\}$ where V is the set of vertices (points) and E is the set of edges. In this case, the problem of matching the points reduces to the classical problem of listing all graph isomorphisms of order 2 (see Appendix A). A graph isomorphism is a permutation of the graph vertices which leaves the graph topologically equivalent. More specifically, given a graph $G = \{E, V\}$, replacing each vertex $i \in V$ with its matched vertex $match(i)$ results in a graph $G' = \{V', E'\}$ such that the set of edges E' equals E. A graph isomorphism of order 2 is an isomorphism where $match(match(i)) = i$ (i.e., either $match(i) = i$, or, $match(i) = j$ and $match(j) = i$). There are several methods for finding all graph isomorphism of order two. We used a simple recursive algorithm for finding these isomorphisms.

4 2D Symmetrization

Dealing with mirror-symmetry and assuming weak perspective projection, a $3D$ mirror-symmetric object has the property that if the projection of the mirror-symmetric pairs of $3D$ points are connected by segments in the $2D$ plane, then all these segments are parallel, i.e., have the same orientation (see Fig. 2). We will denote this property as the "projected mirror-symmetry constraint". If perspective projection is used, these line segments would not be of the same orientation; rather they would be oriented such that the rays extending and including these segments all meet at a single point, which is the epipole [7].

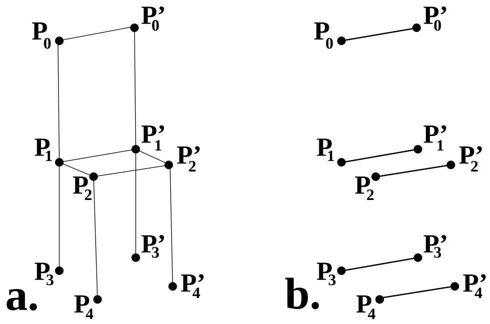

Figure 2: The projected mirror-symmetry constraint. a) A weak perspective projection of a $3D$ mirror-symmetric configuration. Points P_i and P'_i are corresponding mirror-symmetric pairs of points in the $3D$ structure. b) By connecting points P_i with the corresponding P'_i, we obtain a collection of parallel segments.

We use the projected mirror-symmetry constraint to symmetrize the $2D$ data prior to reconstruction of the $3D$ structure. Given a $2D$ configuration of connected points $\{P_i\}_{i=0}^{n-1}$, and given a matching between the points of the configuration (the computation of the matching is described in the previous section), we find a connected configuration of points $\{\hat{P}_i\}_{i=0}^{n-1}$ which satisfy:

1. The configuration of points $\hat{P}_i$ have the same topology as the configuration of points P_i, i.e., points $\hat{P}_i$ and $\hat{P}_j$ are connected if and only if points P_i and P_j are connected.

2. Points $\{\hat{P}_i\}_{i=0}^{n-1}$ satisfy the projected mirror-symmetry constraint, i.e., all the segments connecting points $\hat{P}_i$ and $\hat{P}_j$ (where $\hat{P}_j = match(\hat{P}_i)$) are of the same orientation.

3. The following sum is minimized: $\sum_{i=0}^{n-1} \|P_i - \hat{P}_i\|^2$

It can be shown (see Appendix B) that the points $\{\hat{P}_i\}_{i=0}^{n-1}$ are obtained by projecting each point P_i onto a line at orientation θ passing through the midpoint between P_i and $match(P_i)$, where θ is given by:

$$\tan 2\theta = \frac{2\sum_{i=0}^{n-1}(x_i - match(x_i))(y_i - match(y_i))}{\sum_{i=0}^{n-1}(x_i - match(x_i))^2 - (y_i - match(y_i))^2} \tag{1}$$

Note that two possible solutions exist for Eq. (1). It is easily seen that the solution is achieved when $\sin\theta\cos\theta$ is of opposite sign to the numerator.

Several examples of noisy $2D$ projections of mirror-symmetric configurations of points are shown in Fig. 3 with the closest projected mirror-symmetric configuration, which was obtained using the above algorithm. The matching is shown by the connecting segments.

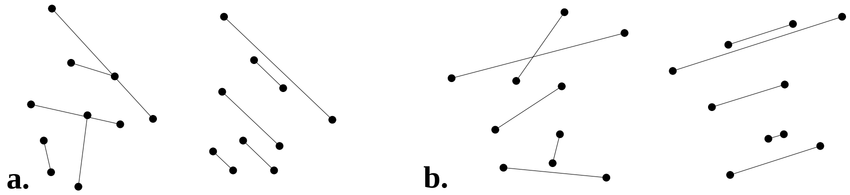

Figure 3: Finding the closest projected mirror-symmetric configuration. a-b) Several examples of noisy $2D$ projections of mirror-symmetric configurations of points (left) and the closest projected mirror-symmetric configuration (right).

5 Experiments

In this section we describe experiments in which $3D$ mirror-symmetric connected configurations are reconstructed from noisy $2D$ perspective projections. We use the two approaches of correction for symmetry which were described in Section 1. The reconstruction method used in the simulations is the invariant reconstruction method described in [16] and reviewed in Appendix C.

The correction procedures were the following:

1. The invariant reconstruction method was applied directly to the $2D$ data with no symmetry assumption. Following the reconstruction, correction for symmetry was applied to the $3D$ reconstruction by finding the closest $3D$ mirror-symmetric configuration using the method described in Section 3.

2. Correction for symmetry was applied to the $2D$ projected data by finding, for every image, the closest projected mirror-symmetric configuration, using the method described in Section 4. Following the correction for symmetry, the reconstruction method was applied to the modified images.

3. Correction for symmetry was performed both prior and following the reconstruction of the $3D$ configuration from $2D$ data.

The reconstruction obtained from these procedures was compared with the original mirror-symmetric $3D$ configuration. The differences were measured by the mean squared-distance between the reconstructed and the original sets of $3D$ points.

5.1 Simulation Results

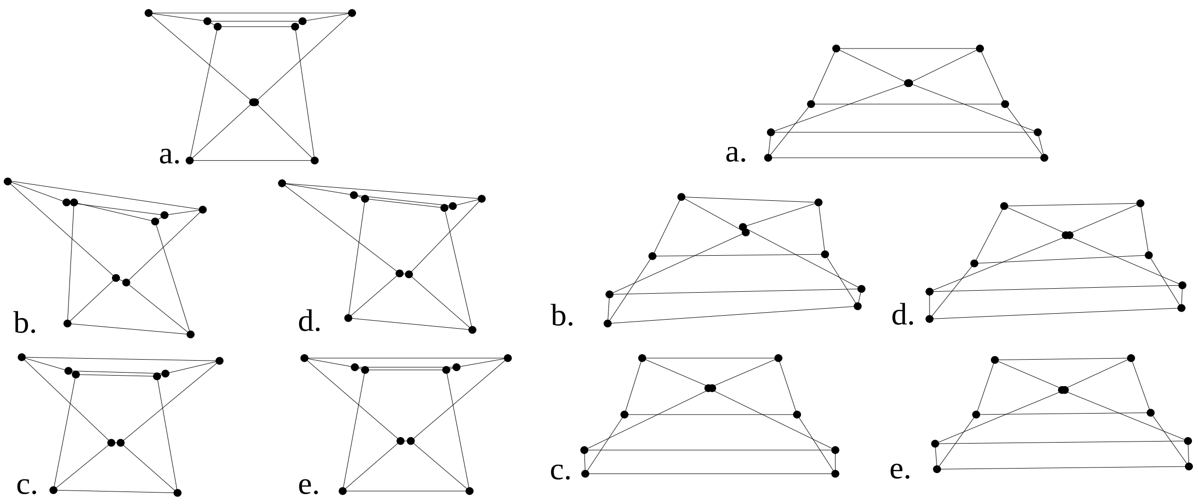

Figure 4: Reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ projections - see text.

Two examples of the simulation are shown in Figure 4. Two randomly chosen $3D$ mirror-symmetric connected configuration of 10 points are shown in Figure 4a. Points were selected randomly in the box $[0, 1]^3$. Eight noisy $2D$ projections were created for each of the $3D$ configurations. Perspective projection was used with a focal length of 5. The projections are from randomly chosen viewpoints and the noise was added to the $2D$ projections and was set at a predefined level of $\sigma = 0.005$ for the first simulation and of $\sigma = 0.05$ for the second simulation.

Reconstruction of the connected configuration directly from the $2D$ projections, with no symmetry assumption, is shown in Figure 4b. The $3D$ reconstruction obtained when correcting for symmetry prior to reconstruction is shown in Figure 4c. The $3D$ reconstruction obtained when correcting for symmetry following the reconstruction is shown in Figure 4d. Finally, Figure 4e shows the $3D$ reconstructed configuration following correction for symmetry prior and following the reconstruction. The differences and percentage of improvement are summarized in Table 1.

In order to obtain some statistical appraisal of the improvement obtained by correcting for symmetry, we applied the simulation many times while varying the simulation parameters. Points were, again, selected randomly in the box $[0, 1]^3$. The number of points was varied between 8 and 24, the number of views was varied between 8 and 24, and the noise level was taken as $\sigma = 0.001, 0.005, 0.01, 0.05$ and 0.1. Every combination of parameters was simulated 300 times. The differences between the reconstruction and the original configuration were measured as in the above two examples.

The percentage of improvement between the reconstruction with no symmetry assumption and the reconstruction with correction for symmetry was calculated and averaged over the

	Sigma	*No Symmetrization*	*Symmetrization prior to reconstruction % improvement*	*Symmetrization following reconstruction % improvement*	*Symmetrization prior & following reconstruction % improvement*
sim. 1	0.005	0.084967	0.072156 15.08%	0.057879 31.88%	0.048645 42.75%
sim. 2	0.05	0.094200	0.086757 7.90%	0.058274 38.14%	0.046645 50.48%

Table 1: The error and % improvement of the reconstructions of $3D$ mirror-symmetric configurations from noisy $2D$ projections.

simulations (7500 trials). The results are given in Table 2. Using σ greater than 0.1 the percentage of improvement breaks down, although when using orthographic projections the improvement is significant up to $\sigma = 0.3$.

σ *(noise)*	*Symmetry prior to reconstruction % improvement*	*Symmetry following reconstruction % improvement*	*Symmetry prior & following reconstruction % improvement*
0.001	11.4	37.7	42.0
0.005	12.6	38.4	43.3
0.01	11.3	38.3	43.2
0.05	4.0	28.9	29.3
0.1	4.8	23.1	22.2
All	8.8	33.3	36.0

Table 2: Improvement in the reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ perspective projections.

5.2 Real data

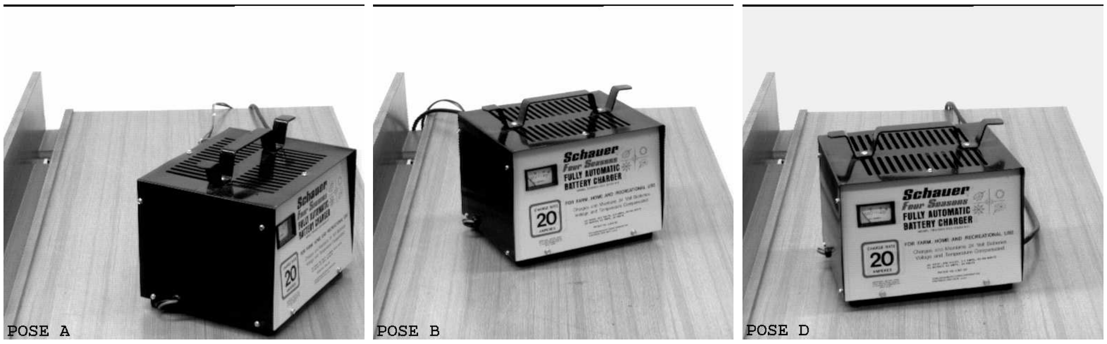

Figure 5: Three $2D$ images of a $3D$ mirror-symmetric object from different view points.

Our algorithm was applied to measurements taken from real $2D$ images of an object. In the first example we took images of the object at three different positions (Fig. 5). 16 feature points were manually extracted from each of the three images. Using the 16 points and the three views, the $3D$ object was reconstructed using the invariant reconstruction method with symmetrization performed prior, following, or both prior and following the reconstruction, as discussed above. The reconstructions were compared to the real (measured) $3D$ coordinates of the object. The results are given in Table 3.

	No Symmetrization	*Symmetrization prior to reconstruction*	*Symmetrization following reconstruction*	*Symmetrization prior & following reconstruction*
error	1.619283	1.388134	1.339260	1.329660
% improvement		14.3	17.3	17.9

Table 3: Improvement in the reconstruction of a real $3D$ mirror-symmetric object from three $2D$ images. The error (average per point) is given in cm, where the object size is aproximately 30cm.

In the second example we took images of the object at five different positions (Fig. 6a). 18 feature points were manually extracted from each of the three images (visually displayed as black crossed in Fig. 6b). The $3D$ object was reconstructed using the invariant reconstruction method with symmetrization performed prior, following, or both prior and following the reconstruction, as discussed above. The reconstructions were compared to the real (measured) $3D$ coordinates of the object. The results are given in Table 4. It can be seen that in this example the symmetrization prior to reconstruction was more effective than following reconstruction. This is due to the fact that the $3D$ reconstruction itself produced a relatively mirror-symmetric object.

a. 

b. 

Figure 6: Reconstruction of a $3D$-mirror-symmetric object from $2D$ images. a. Five $2D$ images of a $3D$ mirror-symmetric chair from different view points. b. The 18 feature points are illustrated by crosses on one of the images.

	No Symmetrization	*Symmetrization prior to reconstruction*	*Symmetrization following reconstruction*	*Symmetrization prior & following reconstruction*
error	3.335983	1.919489	3.192995	1.976036
% improvement		42.5	4.3	40.8

Table 4: Improvement in reconstruction of a real $3D$ mirror-symmetric object from three $2D$ images. The error (average per point) is given in cm, where the object size is aproximately 80cm.

5.3 Discussion

As seen in the examples, the reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ projected data can be greatly improved by correcting for symmetry either prior and/or following reconstruction. Although correcting for symmetry prior to reconstruction improves the result, correcting for symmetry following reconstruction generally gives a greater improvement. Not surprisingly, the greatest improvement in reconstruction is obtained when correction for symmetry is performed both prior and following reconstruction.

Appendix

A The Matching

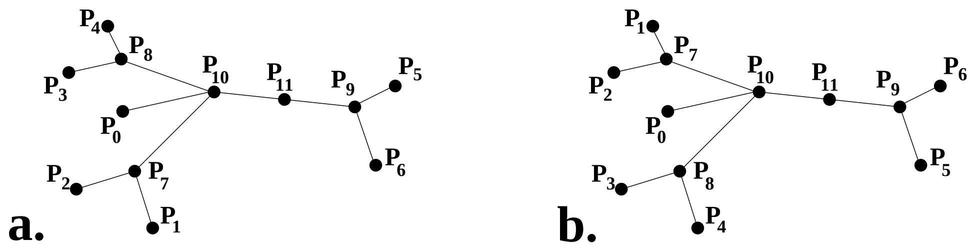

Figure 7: Connected configurations of points. The graph shown in a) is isomorphic to the graph shown in b) (see text).

As described in Section 3 and Section 4, finding the closest mirror-symmetric configuration, or the closest projected mirror-symmetric configuration, requires finding a matching of the points, i.e., a division of the points into sets. Each such set is transformed under the folding/unfolding method into a mirror symmetric set or a projected mirror-symmetric set. Thus, the connectivity of the 2D or 3D points in the original configuration (i.e., the *topology* of the configuration) determines the division of points into sets. As an example, consider the two dimensional connected configuration shown in Fig 7a. Points $P_0, \ldots, P_6$ are leaf nodes and can be paired between them. Points $P_7, \ldots, P_9$ have the same valency of three and can be paired. Points P_{10} and P_{11} stand alone in their valency of 4 and 2 respectively, and will form single-pointpairs (degenerate pairs). Thus a possible division of the points into pairs for measuring mirror-symmetry and for transforming the configuration into a mirror-symmetric configuration is: $\{P_1, P_4\}, \{P_2, P_3\}, \{P_7, P_8\}, \{P_0\}, \{P_{10}\}, \{P_9\}, \{P_5, P_6\}, \{P_{11}\}$.

The valency of a point (the number of edges converging at a point) is necessary but not sufficient in determining the division into sets. Consider the example in Figure 7a. Points P_7 and P_9 have the same valency (3) but obviously cannot be geometrically moved to be mirror-symmetric. This is due to the fact that they are not equivalent in their second order connectivity (i.e. in the valency of their neighboring points): point P_7 has two neighbors of valency 1 and one neighbor of valency 4, whereas point P_9 has two neighbors of valency 1 and one neighbor of valency 2. This reasoning does not stop at the second order connectivity (in Figure 7a points P_3 and P_5 do not agree in their third order connectivity) but must be taken to the maximal connectivity of the configuration (which is equal to the width of the graph).

We consider the original configuration as a graph $G = \{V, E\}$ (as described in Section 3). When considering a configuration of points as a graph, the problem of dividing the points into sets, containing one or two points, reduces to the classical problem of listing all graph isomorphisms of order 2. A graph isomorphism is a permutation of the graph vertices which leaves the graph topologically equivalent. More specifically, given a graph $G = \{E, V\}$, replacing each vertex $i \in V$ with its permuted vertex $match(i)$ results in a graph $G' = \{V', E'\}$ such that the set of edges E' equals E. Note that in this case, for every $(i,j) \in E$ also $(match(i), match(j)) \in E$. A graph isomorphism of order 2 is an isomorphism where $match(match(i)) = i$ (i.e., either $match(i) = i$, or, $match(i) = j$ and $match(j) = i$).

For example, 5 of the 12 isomorphisms of order two for graph $\mathcal{G}$ shown in Figure 7a are:

	$\Pi(P_0)$	$\Pi(P_1)$	$\Pi(P_2)$	$\Pi(P_3)$	$\Pi(P_4)$	$\Pi(P_5)$	$\Pi(P_6)$	$\Pi(P_7)$	$\Pi(P_8)$	$\Pi(P_9)$	$\Pi(P_{10})$	$\Pi(P_{11})$
a	P_0	P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9	P_{10}	P_{11}
b	P_0	P_2	P_1	P_4	P_3	P_6	P_5	P_7	P_8	P_9	P_{10}	P_{11}
c	P_0	P_3	P_4	P_1	P_2	P_5	P_6	P_8	P_7	P_9	P_{10}	P_{11}
d	P_0	P_4	P_3	P_2	P_1	P_5	P_6	P_8	P_7	P_9	P_{10}	P_{11}
e	P_0	P_4	P_3	P_2	P_1	P_6	P_5	P_8	P_7	P_9	P_{10}	P_{11}

Thus by replacing every node with its matched node of isomorphism (e), for example, we obtain the graph G' of Fig 7b which is topologically equivalent to graph G. Notice that every graph isomorphism can be viewed as a permutation of its vertices.

There are several methods for finding all graph isomorphism of order two. We used a simple recursive algorithm for finding this isomorphism.

B Finding the Closest Projected Mirror Symmetry

Given a $2D$ configuration of connected points $\{P_i\}_{i=0}^{n-1}$, and given a matching between the points of the configuration (i.e., $\forall$ P_i define $match(P_i)$ where $match(match(P_i)) = P_i$ and where $match(P_i) = P_i$ is permissible), we find a connected configuration of points $\{\hat{P}_i\}_{i=0}^{n-1}$ which satisfy:

1. The configuration of points $\hat{P}_i$ have the same topology as the configuration of points P_i, i.e., points $\hat{P}_i$ and $\hat{P}_j$ are connected if and only if points P_i and P_j are connected.
2. Points $\{\hat{P}_i\}_{i=0}^{n-1}$ satisfy the projected mirror-symmetry constraint, i.e., all the segments connecting points $\hat{P}_i$ and $\hat{P}_j$ (where $\hat{P}_j = match(P_i)$) are of the same orientation.
3. The following sum is minimized:

$$\sum_{i=0}^{n-1} \|P_i - \hat{P}_i\|^2 \tag{2}$$

Consider first a simple case where we are given two points P_0 and P_1 in $\mathcal{R}^2$ and an orientation θ (without loss of generality, θ is the angle to the positive x-axis). We find 2 points $\hat{P}_0$ and $\hat{P}_1$ such that the segment connecting them is at orientation θ and the following sum is minimized:

$$||P_0 - \hat{P}_0||^2 + ||P_1 - \hat{P}_1||^2 \tag{3}$$

Claim 1: *Given a line $y = \tan(\theta)x + c$ ($c \in \mathcal{R}$), points $\hat{P}_0$ and $\hat{P}_1$ which minimize Eq. (3) are obtained by projecting P_0 and P_1 respectively onto the line (Fig. 8a).*

Furthermore, the line of orientation θ, on which positioning points $\hat{P}_0$ and $\hat{P}_1$ minimizes Eq. (3), passes through the centroid (or mid-section point) of P_0 and P_1 (Fig. 8b).

The claim follows from the fact that the projection of a point on a line is the closest location on the line to the point.

Thus, given 2 points $P_0, P_1 \in \mathcal{R}^2$ and given an orientation θ, the value of Eq. (3) is:

$$||P_0 - \hat{P}_0||^2 + ||P_1 - \hat{P}_1||^2 = \frac{1}{2}[(x_1 - x_0)\sin(\theta) - (y_1 - y_0)\cos(\theta)]^2$$

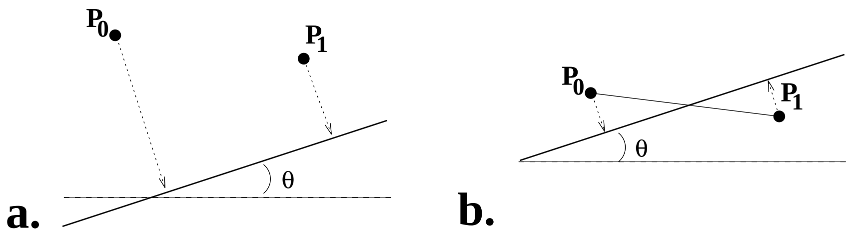

Figure 8: Finding the closest projected mirror-symmetric configuration, a simple case of two points: a) given two points P_0 and P_1 and given a line $y = \tan(\theta) + c$ ($c \in \mathcal{R}$), the points closest to P_0 and P_1 which lie on the line are obtained by projection; b) the points lying on a line of orientation θ that are closest to P_0 and P_1 are obtained by projecting P_0 and P_1 onto a line of orientation θ passing through the midpoint between P_0 and P_1.

Consider now n $2D$ points $\{P_i\}_{i=0}^{n-1}$ and a given matching of these points. In order to find the points $\{\hat{P}_i\}_{i=0}^{n-1}$ that minimize Eq. (2) and that satisfy the projected mirror-symmetry constraint, we must find the orientation θ which minimizes Eq. (2). For a given orientation θ, the value of Eq. (2) is

$$\sum_{i=0}^{n-1} ||P_i - \hat{P}_i||^2 = \sum_{i=0}^{n-1} [(x_i - match(x_i))\sin(\theta) - (y_i - match(y_i))\cos(\theta)]^2$$

where $match(x_i)$,$match(y_i)$ are the $2D$ coordinates of the point $match(P_i)$.

Taking the derivative with respect to θ and equating to zero we obtain for the minimal θ:

$$\tan 2\theta = \frac{2\sum_{i=0}^{n-1}(x_i - match(x_i))(y_i - match(y_i))}{\sum_{i=0}^{n-1}(x_i - match(x_i))^2 - (y_i - match(y_i))^2} \tag{4}$$

As noted in Section 4, two possible solutions exist for Eq. (4). It is easily seen that the minimum is achieved when $\sin\theta\cos\theta$ is of opposite sign to the numerator of Eq. (4).

Thus we have a closed form solution for finding the closest projected mirror-symmetric set of points; given the $2D$ points $\{P_i\}_{i=0}^{n-1}$ and a matching $\{match(P_i)\}_{i=0}^{n-1}$:

1. calculate the optimal orientation θ using Eq. (1).
2. calculate the coordinates of each point $\hat{P}_i$ by projecting the points P_i onto a line at orientation θ passing through the midpoint between P_i and $match(P_i)$.

C Review of the invariant reconstruction algorithm

This linear method was described in [16]. It computes an invariant description of the Euclidean structure of points from a sequence of images assuming weak perspective.

Let $\{\mathbf{p}_i\}_{i=0}^{n-1}$, $\mathbf{p}_i \in \mathcal{R}^3$, denote the $3D$ coordinates of an object composed of n features in some Cartesian coordinate system. For simplicity and clarity, we start with the case $n = 4$ and $\mathbf{p}_0 = (0,0,0)$. Let P denote the 3×3 matrix whose columns are the vectors $\{\mathbf{p}_i\}_{i=1}^{n-1}$, namely, $P = [\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3]$.

A representation of the object shape, which is invariant to rigid transformations of the camera, is the Gramian matrix[1] $G = P^T P$:

$$G = \begin{pmatrix} \mathbf{p}_1^T\mathbf{p}_1 & \mathbf{p}_1^T\mathbf{p}_2 & \mathbf{p}_1^T\mathbf{p}_3 \\ \mathbf{p}_1^T\mathbf{p}_2 & \mathbf{p}_2^T\mathbf{p}_2 & \mathbf{p}_2^T\mathbf{p}_3 \\ \mathbf{p}_1^T\mathbf{p}_3 & \mathbf{p}_2^T\mathbf{p}_3 & \mathbf{p}_3^T\mathbf{p}_3 \end{pmatrix}$$

Using the weak pespective approximation, it can be shown [15] that:

$$\begin{aligned} \mathbf{x}^T G^{-1}\mathbf{x} &= \mathbf{y}^T G^{-1}\mathbf{y} \\ \mathbf{x}^T G^{-1}\mathbf{y} &= 0 \end{aligned} \tag{5}$$

where the vectors $\mathbf{x} = (x_1, x_2, x_3)$ and $\mathbf{y} = (y_1, y_2, y_3)$ are obtained from the image data points $\{P_i = (x_i, y_i)\}_{i=1}^3$. We compute the Gramian of the 4 points by solving the linear system of equations given in Eq. (5) (note that Eq. (5) is linear in the elements of the inverse Gramian). The Gramian gives the complete Euclidean-invariant (metric) structure of the 4 points [15].

Given more than 4 points, the algorithm proceeds as follows:

- select 4 basis points from the data (using QR factorization to maximize the independence of the selected points);
- compute the affine structure of all the points by solving a linear system of equations;

[1] computing depth (matrix P) from G is straightforward, and computationally very fast (a decomposition known as **Choleski** factorization).

- compute the Euclidean structure of the 4 basis points by solving a linear system (given in Eq. (5));
- obtain the Euclidean structure of all the points if necessary (this can be done by multiplying a vector of affine coordinates by the root of the Gramian G of the basis points).

References

[1] K.S. Arun, T.S. Huang, and S.D. Blostein. Least squares fitting of two 3D point sets. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 9(5):698–700, Sept. 1987.

[2] J. Bigün. Recognition of local symmetries in gray value images by harmonic functions. In *International Conference on Pattern Recognition*, pages 345–347, 1988.

[3] A. Blake, M. Taylor, and A. Cox. Grasping visual symmetry. In *International Conference on Pattern Recognition*, pages 724–733, Berlin, May 1993.

[4] M. Brady and H. Asada. Smoothed local symmetries and their implementation. *International Journal of Robotics Research*, 3(3):36–61, 1984.

[5] S. Friedberg and C. Brown. Finding axes of skewed symmetry. *Computer Vision, Graphics, and Image Processing*, 34:138–155, 1986.

[6] M-K. Hu. Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, IT-20:179–187, Feb 1962.

[7] H. Mitsumoto, S. Tamura, K. Okazaki, N. Kajimi, and Y. Fukui. 3-d reconstruction using mirror images based on a plane symmetry recovering method. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(9):941–946, 1992.

[8] W.G. Oh, M. Asada, and S. Tsuji. Model based matching using skewed symmetry information. In *International Conference on Pattern Recognition*, pages 1043–1045, 1988.

[9] T. Poggio and T. Vetter. Recognition and structure from one 2D model view: Observations on prototypes, object classes and symmetries. Technical Report A.I. Memo No. 1347, MIT, February 1992.

[10] J. Ponce. On characterizing ribbons and finding skewed symmetries. *Computer Vision, Graphics, and Image Processing*, 52:328–340, 1990.

[11] D. Reisfeld, H. Wolfson, and Y. Yeshurun. Robust facial feature detection using local symmetry. In *International Conference on Pattern Recognition*, pages A:117–120, Champaign, USA, June 1990.

[12] C.A. Rothwell, D.A. Forsyth, A. Zisserman, and J.L. Mundy. Extracting projective structure from single perspective views of 3D point sets. In *International Conference on Pattern Recognition*, pages 573–582, Berlin, May 1993.

[13] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry seeking models and object reconstruction. *Int. J. Computer Vision*, 1:211–221, 1987.

[14] F. Ulupinar and R. Nevatia. Using symmetries for analysis of shape from contour. In *International Conference on Computer Vision*, pages 414–426, 1988.

[15] D. Weinshall. Model-based invariants for 3D vision. *International Journal of Computer Vision*, 10(1):27–42, 1993.

[16] D. Weinshall and C. Tomasi. Linear and incremental acquisition of invariant shape models from image sequences. In *International Conference on Computer Vision*, pages 675–682, Berlin, Germany, 1993.

[17] H. Zabrodsky. *Computational Aspects of Pattern Characterization - Continuous Symmetry*. PhD thesis, Hebrew University, Jerusalem, Israel, 1993.

[18] H. Zabrodsky, S. Peleg, and D. Avnir. Hierarchical symmetry. In *International Conference on Pattern Recognition*, volume C: Image, Speech, and Signal Analysis, pages 9–12, The Hague, August-September 1992.

[19] H. Zabrodsky, S. Peleg, and D. Avnir. A measure of symmetry based on shape similarity. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 703–706, Champaign, June 1992.