Using Bilateral Symmetry To Improve 3D Reconstruction From Image Sequences

Hagit Zabrodsky and Daphna Weinshall
Institute of Computer Science
The Hebrew University of Jerusalem
91904 Jerusalem, Israel
contact email: daphna@cs.huji.ac.il

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 demonstrate how bilateral symmetry can be used to improve the accuracy in $3D$ reconstruction. The symmetry property is used 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.

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, 18] 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 methods (or structure from a sequence of $2D$ images), 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). As a working example, 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.

Exploiting the fact that the $3D$ structure to be reconstructed is mirror-symmetric, we incorporate a symmetrization procedure into the reconstruction scheme as a separate module, independent of the reconstruction method used. We employ two approaches:

- correct for bilateral-symmetry prior to reconstruction

- correct for bilateral-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 3.1 for finding the closest mirror-symmetric configuration to a given $3D$ connected configuration.

Correction for symmetry prior to reconstruction requires application of some symmetrization procedure to the $2D$ data with respect to the $3D$ symmetry. In Section 3.2 we describe a symmetrization procedure of $2D$ data for projected $3D$ mirror-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 outline the reconstruction algorithms, which start off by verifying that the object is symmetrical, and proceed by using one of the different symmetrization methods described above to improve the reconstruction. 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 [2, 6] 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 [5, 8]. A method for dealing with any type of symmetry in both $2D$ and $3D$ has recently been suggested [17, 18].

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 [7, 9, 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 objects, which is the subject of the present paper.

3 Symmetrization algorithms and their complexity

In Section 3.1 we describe a closed-form method for finding the closest mirror-symmetric configuration to a given $3D$ connected configuration. In Section 3.2 we describe a symmetrization procedure of $2D$ data for projected $3D$ mirror-symmetry. A complexity analysis of both symmetrization procedures is given in Section 3.3.

3.1 3D symmetrization

In [17, 18] 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 [17, 18]. Below we describe another *closed-form* solution that gives equivalent results in the case of $3D$ mirror-symmetry.

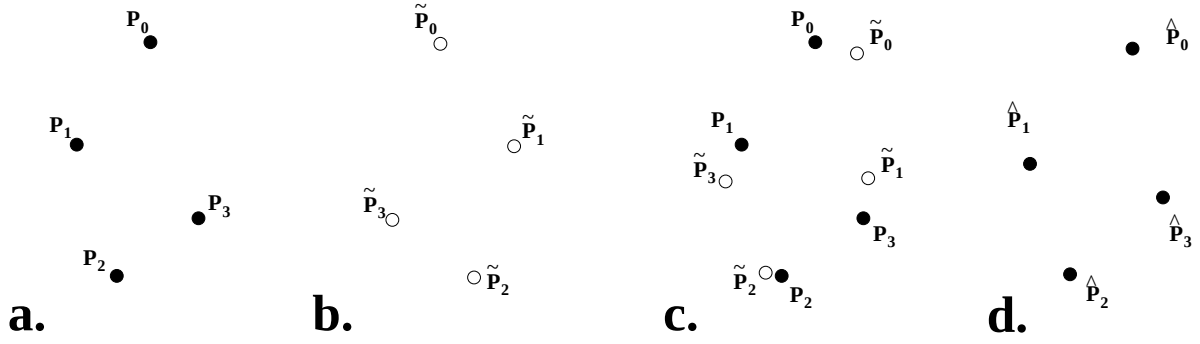

Figure 1: Obtaining the closest mirror symmetric set of points - see text.

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).

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:

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

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 Distance 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 if connectivity or ordering information is available regarding the configuration of points. In some cases the complexity is reduced from exponential to linear (see below, Section 3.3).

3.2 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 [7].

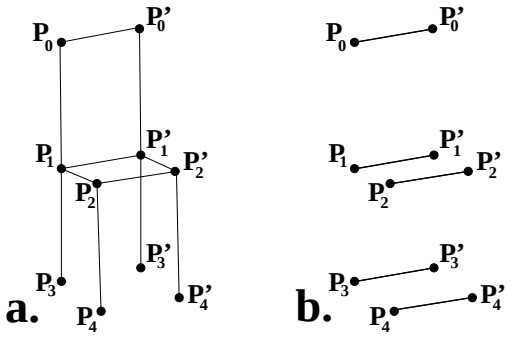

Figure 2: The projected mirror-symmetry constraint. a) A weak perspective projection of a $3D$ mirror-symmetric configuration. Points P_i and P_i' are corresponding mirror-symmetric pairs of points in the $3D$ structure. b) By connecting points P_i with the corresponding P_i', we obtain a collection 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 (see Section 3.3 for discussion on matching), 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 lines passing through points $\hat{P}_i$ and $\hat{P}_j$ (where $\hat{P}_j = match(\hat{P}_i)$) are of the same orientation.

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

It can be shown (see Appendix A) 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:

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

Note that two possible solutions exist for Eq. (2). 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.

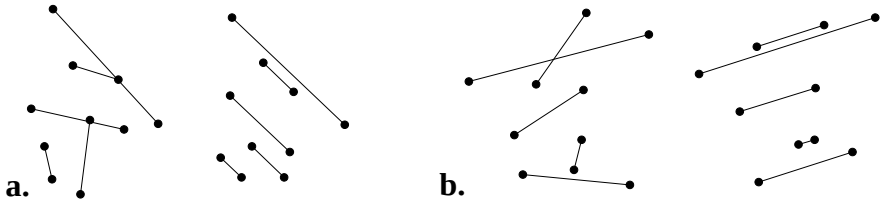

Figure 3: Finding the closest projected mirror-symmetric configuration. a-b) Several examples of noisy $2D$ projections of mirror-symmetric configurations of points (left) and the closest projected mirror-symmetric configuration (right).

3.3 A Complexity analysis: matching and symmetrization

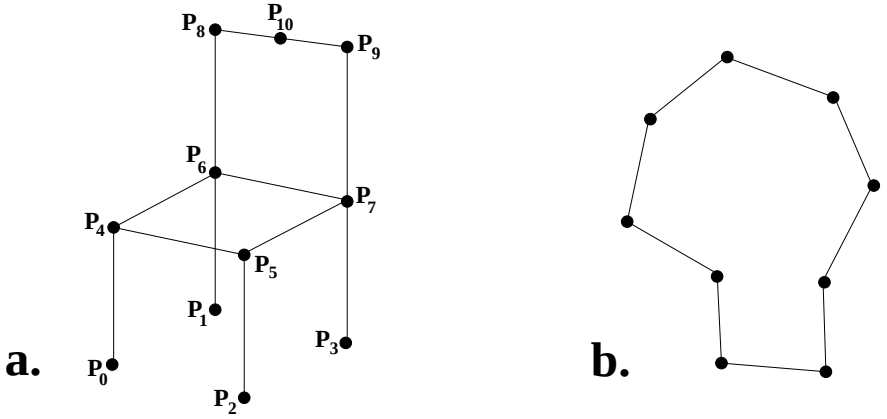

Figure 4: Connected configurations of points. The connectivity in the graph constrain the possible matchings of the points (see text).

In terms of complexity, the crucial step in the symmetrization algorithms described above is the minimization of the Symmetry Distance value over all possible matchings of feature points. Offhand, matching of feature points is of exponential complexity. However, as will be discussed below, the actual computational costs can be greatly reduced by constraining the search space of all possible matches.

Graph matching:

As described in Sections 3.1,3.2, 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. We assume that the folding/unfolding method maintains any connectivity and any ordering that exists between the points of the configuration. Thus, the connectivity of the 2D or 3D points in the original configuration (i.e., the *topology* of the configuration) constrains the division of points into sets.

For example, consider the connected configuration shown in Fig 4a. Points $P_0, \ldots, P_3$ are leaf nodes and can only be paired between themselves. Points P_6 and P_7 are the only nodes with valency of 4 (where valency denotes the number of edges converging at a point), and thus must be paired between themselves or form single-point pairs. Point P_{10} stands alone in its valency of 2 and can only form a single-point pair (a degenerate pair).

Further constraints on possible matching of points are obtained when taking into consideration that the valency of a point is necessary but not sufficient in determining the division into sets. In the example of Figure 4a, points P_2 and P_3 have the same valency (1) 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_2 has a neighbor of valency 3, whereas point P_3 has a neighbor of valency 4. This reasoning does not stop at the second order connectivity but must be taken to the maximal connectivity of the configuration (which is equal to the width of the graph).

These considerations constrain the space of all possible matchings that can give rise to a mirror-symmetric configuration. Specifically, for the example of Figure 4a, the number of possible matchings reduces to 2, namely: $\{P_0, P_2\}, \{P_1, P_3\}, \{P_4, P_5\}, \{P_6, P_7\}, \{P_8, P_9\}, \{P_{10}\}$ and the matching in which all pairs are degenerate pairs. For the class of cyclically connected configurations (as that shown in Figure 4b) it can be shown that the number of possible matchings is reduced form exponential to linear (specifically, for a cyclic configuration of n points there are n possible matchings).

More generally, we consider the original configuration as a graph $G = \{V, E\}$. 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 $match(j) = i$). The constraints discussed above are inherent in any algorithm that finds graph isomorphisms. There are several methods for finding all graph isomorphism of order two. We used a simple recursive algorithm for finding this isomorphism.

Points matching:

In some cases, it is difficult to extract connectivity and order information relating to the original configuration of points. In these cases the number of possible matchings increases exponentially with the number of points in the configuration. For these cases a heuristic approach can be used instead. The above described approach of graph isomorphism assumes a matching is to be found

prior to finding the optimal reflection plane. We now consider the problems of point matching and of finding the optimal reflection plane as confounded; given a matching, we can determine the optimal reflection plane (using the folding/unfolding algorithm described in Section 3.1), and given the reflection plane we can constrain the possible matchings and more easily determine the pairings.

We suggest a heuristic approach which tries to solve these two problems simultaneously as follows. For every possible pair of points we determine the corresponding reflection plane (the plane perpendicular to and passing through the mid point of the segment connecting the two points). We build a histogram of all possible reflection planes for pairs from the original set of points. Peaks in the histogram will point at candidates for the optimal reflection plane. Given these reflection planes, the matchings can be determined and the Symmetry Distance evaluated.

4 Algorithms

We propose reconstruction algorithms that enhance existing structure from motion algorithms by using symmetry to improve the reconstruction. Two features characterize our approach:

- The underlying method is independent of the particular reconstruction algorithm, i.e., the enhancement stage can be used together with any reconstruction algorithm. In the examples below we use the particular reconstruction algorithm described in Appendix B. We believe, however, that the improvement obtained is general, and characterizes the properties of the enhancement stage independently of the particular reconstruction method used.
- An enhanced algorithm should only be used in the reconstruction of mirror symmetrical objects. Thus we develop a test to measure object symmetry from the given data. If the object is not likely to be symmetrical, the enhancement stage is not applied - reconstruction is done with the bare (un-enhanced) reconstruction algorithm.

More specifically, our enhanced reconstruction approach is the following:

1. **Pre-processing:**
   (a) Select a reconstruction algorithm.
   (b) Test whether the object is bilaterally symmetrical using the 2D Symmetry Distance defined in Section 3.2. This measure can be used to determine whether the configuration of $2D$ points is indeed a projection of a $3D$ mirror-symmetric configuration. Specifically, under the condition that the system noise is bounded, if the Symmetry Distance is large, we may assume that we are not dealing with a $3D$ symmetric configuration. If the Symmetry Distance is small in all projections (in all images in the given sequence), we may assume that the $3D$ configuration is symmetric and that any deviations are due to noise.
2. **Symmetry enhanced reconstruction:** if the symmetry test is successful, we propose three algorithms for improving $3D$ reconstruction from noisy $2D$ perspective projections (image sequences) using symmetry.

(a) The $3D$ reconstruction method is applied directly to the $2D$ data with no symmetry assumption. Following the reconstruction, correction for symmetry is applied to the $3D$ reconstruction by finding the closest $3D$ mirror-symmetric configuration using the method described in Section 3.1.

(b) Correction for symmetry is applied to the $2D$ projected data by finding, for every image, the closest projected mirror-symmetric configuration, using the method described in Section 3.2. Following the correction for symmetry, the reconstruction method is applied to the modified images.

(c) Correction for symmetry is performed both prior and following the reconstruction of the $3D$ configuration from $2D$ data.

5 Experiments

In this section we describe experiments in which the three algorithms described in Section 4 are compared and evaluated. As a working example, we demonstrate and compare the algorithms using the invariant reconstruction method described in [16] and reviewed in Appendix B.

The reconstruction was performed on both simulated and real data. The reconstruction obtained from the three 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

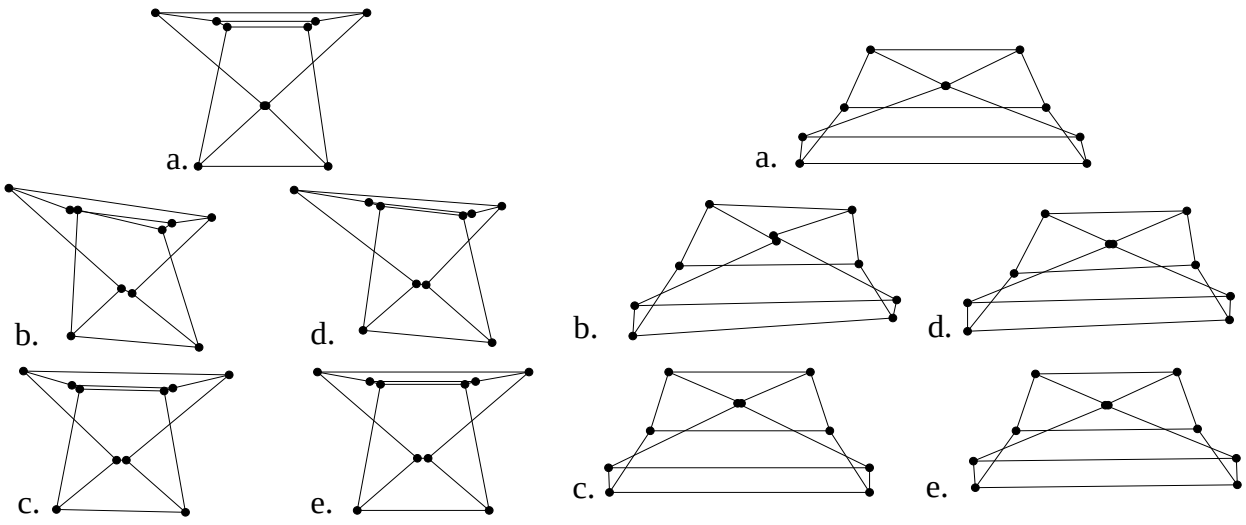

Figure 5: Reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ projections - see text.

Two examples of the simulation are shown in Figure 5. Two randomly chosen $3D$ mirror-symmetric connected configuration of 10 points are shown in Figure 5a. 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 5b. The $3D$ reconstruction obtained when correcting for symmetry prior to reconstruction is shown in Figure 5c. The $3D$ reconstruction obtained when correcting for symmetry following the reconstruction is shown in Figure 5d. Finally, Figure 5e 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.

	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 reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ projections.

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 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 $\sigma = 0.3$.

5.2 Real data

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. 6). 16 feature points were manually extracted from each of the three images. The points were automatically matched using the heuristical method described in Section 3.3, giving 8 pairs of symmetrical points. 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

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.

Figure 6: Three $2D$ images of a $3D$ mirror-symmetric object from different view points.

discussed above. The reconstructions were compared to the real (measured) $3D$ coordinates of the object. The results are given in Table 3.

In the second example we took images of the object at five different positions (Fig. 7a). 18 feature points were manually extracted from each of the three images (visually displayed as black crosses in Fig. 7b). 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.

6 Discussion

The work described above shows that existing reconstruction methods can be enhanced, and their output significantly improved, if bilateral symmetry constraints are used during the reconstruction of symmetrical objects. As seen in the examples above, the reconstruction of $3D$ mirror-symmetric

	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 approximately 30cm.

a.

b.

Figure 7: 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 approximately 80cm.

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.

We believe this improvement to be independent of the reconstruction method that had been used. In particular, our examples should not be taken as a comparison between different reconstruction methods. Rather, we demonstrate the improvement that can **always** be obtained when using valid symmetry constraints. The only valid comparison is between the different ways of applying these symmetry constraints, which give rise to different methods of enhancing **any** reconstruction method.

Appendix

A 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 lines passing through 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:
$$\sum_{i=0}^{n-1} \|P_i - \hat{P}_i\|^2 \tag{3}$$

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:

$$\|P_0 - \hat{P}_0\|^2 + \|P_1 - \hat{P}_1\|^2 \tag{4}$$

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. (4) 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. (4), 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. (4) is:

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

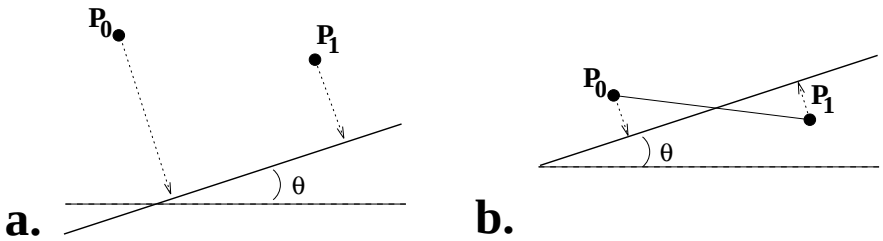

Figure 8: Finding the closest projected mirror-symmetric configuration, a simple case of two points: a) given two points P_0 and P_1 and given a line $y = \tan(\theta) + c$ ($c \in \mathcal{R}$), the points closest to P_0 and P_1 which lie on the line are obtained by projection; b) the points lying on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

where (x_i, y_i) are the coordinates of point P_i.

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. (3) and that satisfy the projected mirror-symmetry constraint, we must find the orientation θ which minimizes Eq. (3). For a given orientation θ, the value of Eq. (3) 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:

$$\tan 2\theta = \frac{2\sum_{i=0}^{n-1}(x_i - match(x_i))(y_i - match(y_i))}{\sum_{i=0}^{n-1}(x_i - match(x_i))^2 - (y_i - match(y_i))^2} \tag{5}$$

As noted in Section 3.2, two possible solutions exist for Eq. (5). It is easily seen that the minimum is achieved when $\sin\theta\cos\theta$ is of opposite sign to the numerator of Eq. (5).

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. (5).

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 $match(P_i)$.

B 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 perspective approximation, it can be shown [15] that:

$$\begin{aligned} \mathbf{x}^T G^{-1} \mathbf{x} &= \mathbf{y}^T G^{-1} \mathbf{y} \\ \mathbf{x}^T G^{-1} \mathbf{y} &= 0 \end{aligned} \tag{6}$$

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. (6) (note that Eq. (6) 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;
- compute the Euclidean structure of the 4 basis points by solving a linear system (given in Eq. (6));
- 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 Computer Vision*, 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.

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

[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 detection of facial features by generalized symmetry. In *International Conference on Pattern Recognition*, pages A:117–120, June 1992.

[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 Computer Vision*, 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. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(5):512–517, 1995.

[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. Symmetry as a continuous feature. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(12):1154–1166, 1995.