ECCV-94 conf. proc., p.403-410, May 1994

Utilizing symmetry in the reconstruction of three-dimensional shape from noisy images

Hagit Zabrodsky[1] and Daphna Weinshall[1]

Institute of Computer Science
The Hebrew University of Jerusalem
91904 Jerusalem, Israel
contact email: gigi@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 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.

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 [4] 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.

The reconstruction of a general $3D$ structure from $2D$ projections, or the problem of structure from motion, is widely studied in computer vision and many reconstruction algorithms have been proposed. In this work 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 [3] with the method dealing with inexact symmetries suggested in [4], for improving the input and output data in the structure reconstruction from several views. Previous work on exploiting symmetry is described in [4].

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 2 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 4, the error in reconstruction is greatly reduced. In Section 3 we describe a symmetrization procedure of $2D$ data for projected $3D$ mirror-symmetry. In Section 4 we give examples and comparisons between correction for symmetry prior and following $3D$ reconstruction, using real and simulated data.

2 3D symmetrization

In [5, 4] 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 [5, 4]. Below we describe a *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):

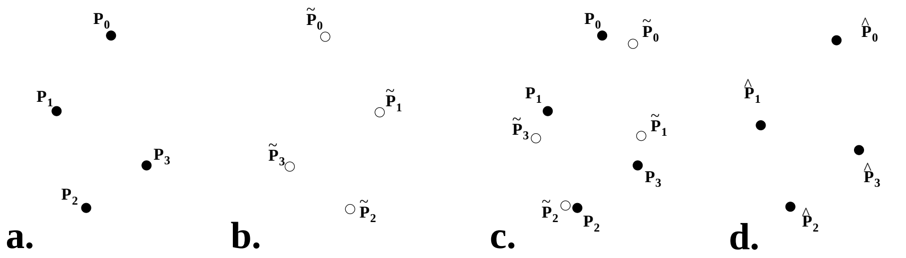

Fig. 1. Obtaining the closest mirror symmetric set of points - see text.

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. **Average** each original point P_i with its reflected matched point P_j, obtaining the point $\hat{P}_i$ (Fig. 1d). The points $\{\hat{P}_i\}_{i=0}^{n-1}$ are mirror symmetric.
5. **Evaluate** the symmetry distance value: $\frac{1}{n}\sum_{i=0}^{n-1} ||P_i - \hat{P}_i||^2$.
6. **Minimize** the symmetry distance value obtained in Step 5 by repeating Steps 1-5 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 [4]).

In practice, the minimization in Step 6 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. 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.

3 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 [2].

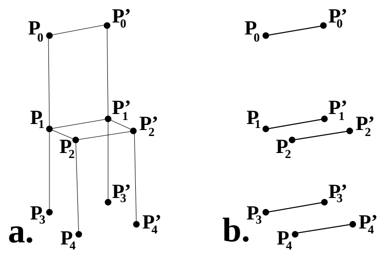

Fig. 2. The projected mirror-symmetry constraint. a) A weak perspective projection of a $3D$ mirror-symmetric configuration. Points P_i and P_i' are corresponding mirror-symmetric pairs of points in the $3D$ structure. b) By connecting points P_i with the corresponding P_i', we obtain a collection of 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 $\hat{P}_j = match(P_i)$) 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 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.

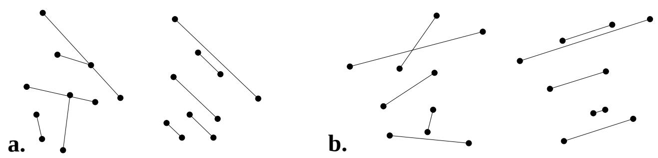

Fig. 3. Finding the closest projected mirror-symmetric configuration. a-b) Several examples of noisy $2D$ projections of mirror-symmetric configurations of points (left) and the closest projected mirror-symmetric configuration (right).

4 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 [3].

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 2.
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 3. 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.

4.1 Simulation Results

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.

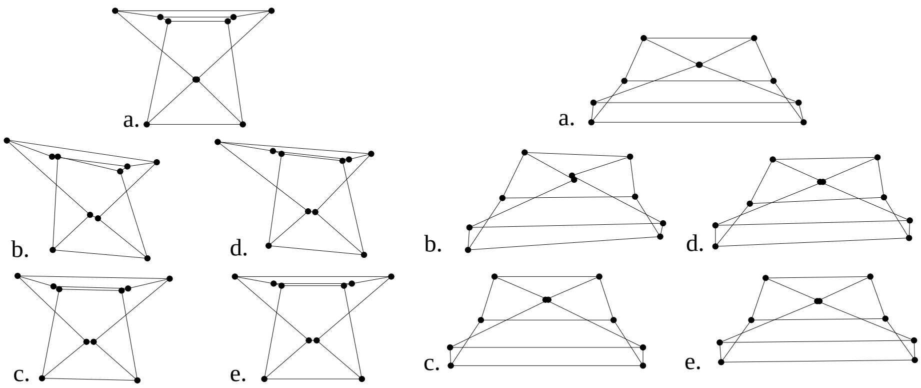

Fig. 4. Reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ projections - see text.

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.

	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.

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 $\sigma = 0.3$.

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

4.2 Real data

Our algorithm was applied to measurements taken from $2D$ real images of an object. In the following 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.

4.3 Discussion

As seen in the examples, reconstruction of $3D$ mirror-symmetric configurations from noisy $2D$ projected data can be greatly improved by correcting for symme-

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

	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 reconstruction of a real $3D$ mirror-symmetric object from three $2D$ images.

try 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.

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. 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.
3. 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.
4. H. Zabrodsky. *Computational Aspects of Pattern Characterization - Continuous Symmetry*. PhD thesis, Hebrew University, Jerusalem, Israel, 1993.
5. 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.

This article was processed using the LaTeX macro package with LLNCS style