Symmetric 3D objects are an easy case for 2D object recognition

Thomas Vetter and Tomaso Poggio

Symmetry and object recognition

Publised in:
Spatial Vision 8:4 (1994) 443-453.

Abstract

According to the **1.5 views theorem** (Ullman and Basri, 1991; Poggio, 1990) recognition of a specific 3D object (defined in terms of pointwise features) from a novel 2D view can be achieved from at least two 2D model views (for each object, for orthographic projection). In this note we discuss how recognition can be achieved from a single 2D model view by exploiting prior knowledge of an object's symmetry. We prove that for any bilaterally symmetric 3D object one non-accidental 2D model view is sufficient for recognition since it can be used to generate additional "virtual" views. We also prove that for bilaterally symmetric objects the correspondence of four points between two views determines the correspondence of **all** other points. Symmetries of higher order allow the recovery of Euclidean structure from a single 2D view.[1]

1 Introduction

Image-based techniques for object recognition have recently been developed to recognize a specific three-dimensional object after a "learning" stage, in which a few two-dimensional views of the object are used as training examples (Poggio and Edelman, 1990; Edelman and Poggio, 1992). A theoretical lower bound on the number of views is provided by the *1.5-views theorem* (Poggio, 1990; Ullman and Basri, 1991; for more details see section 2.1 in this paper). In the orthographic case, this theorem implies that two views – defined in terms of pointwise features – are sufficient for recognition or equivalent to define the affine structure of an object (see also Koenderink and van Doorn, 1991). It is known that in the case of perspective projection, two views are sufficient to compute projective invariants specific to the object (Faugeras, 1992; Hartley et al.1992 and Shashua, 1993). Under more general conditions (more general definition of "view", non-uniform transformations etc.) and, depending on the implementation, many more views may be required (Poggio and Edelman's estimate is on the order of 100 for the whole viewing sphere using their approximation network).

Though this is an easily-satisfied requirement in many cases, there are situations in which only one 2D view is available as a model. As an example, consider the problem of recognizing a face from just one view: humans can do it, even for different facial expressions (of course an almost-frontal view may not be sufficient for recognizing a

[1] Part of this paper has appeared as MIT AI Lab. Memo No. 1347.

profile view and in fact the praxis of person-identification requires usually a frontal *and* a side view).

Clearly one single view of a generic 3D object (if shading is neglected) does not contain sufficient 3D information. If, however, the object belongs to a class of similar objects, it seems possible to infer appropriate transformations for the class and use them to generate other views of the specific object from just one 2D view of it. We certainly are able to recognize faces which are slightly rotated from just one quasi-frontal view, presumably because we exploit our extensive knowledge of the typical 3D structure of faces.

One can pose the following problem: *is it possible from one 2D view of a 3D object to generate other views of that object, exploiting knowledge of the legal transformations associated with objects of the same class?* (We call a 2D transformation of a 2D view *legal* if its result is identical to the projection onto the image-plane of a rigid rotation of the unknown 3D object.) A positive answer would imply (for orthographic projection, uniform affine transformations and in the absence of self-occlusions) that a novel 2D view may be recognized from a single 2D model view, because of the 1.5-views theorem.

In this paper we consider the case in which legal transformations for a specific object (i.e. transformations that generate new correct views from a given one) immediately are available as a property of the class. In particular, we will discuss certain symmetry properties.

The main results of the paper are two.

1. We prove that, for any bilaterally symmetric 3D object (such as a face), one 2D model view is sufficient for recognition of a novel 2D view (for orthographic projection and uniform affine transformations). This result is equivalent to the following statement: for bilaterally symmetric objects, a model-based recognition invariant (as defined by Weinshall, 1993) can be learned from just one model 2D view. It is also closely related to the projective invariant computed on symmetric objects by Rothwell et al. 1993.

2. We also prove that for symmetries of higher order (such as two-fold symmetries, i.e., bilateral symmetry with respect to two symmetry planes) it is possible to recover Euclidean structure from one 2D view (see also Kontsevich, 1993).

In the final section, we briefly mention some of the implications of our results for the practical recognition of bilaterally-symmetric objects, for human perception of 3D structure from single views of geometric objects and, more generally, for the role of symmetry detection in human vision.

2 Results

2.1 Recognition from One 2D Model View

2.1.1 Generating 'virtual' views

Suppose that we have a single 2D model view of a 3D object, which is defined in terms of pointwise features. A 2D view can be represented by a vector $X = (x_1, y_1, x_2, y_2 ......x_n, y_n)$, that is by the x, y-coordinates of its n feature points. Assume further that (a) we know *a priori* that the object is bilaterally symmetric (for instance, because we identify the class to which it belongs and we know that this class has the property of bilateral symmetry) and (b) we find in the 2D view the correspondence of the symmetric pairs of points. It can be shown that for views of *bilaterally symmetric* objects there exist 2D transformations D on a pair p of symmetric points of the object that yield a *legal view* p^* of the pair. This new view is the projection of a rigid rotation of the unknown 3D object onto the image-plane

$$Dp_{pair} = p^*_{pair}. \tag{1}$$

Under the transformations D_1, D_2, and D_3

$$D_1 = \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \\ -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \quad D_2 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \quad D_3 = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

the pair of symmetric points p_0 transforms into p_1, p_2 and p_3

$$\mathbf{p}_0 = \begin{pmatrix} x_R \\ y_R \\ x_L \\ y_L \end{pmatrix}, \quad p_1 = \begin{pmatrix} -x_L \\ y_L \\ -x_R \\ y_R \end{pmatrix}, \quad p_2 = \begin{pmatrix} x_L \\ y_L \\ x_R \\ y_R \end{pmatrix}, \quad p_3 = \begin{pmatrix} -x_R \\ y_R \\ -x_L \\ y_L \end{pmatrix}.$$

Each of these transformations applied to all symmetric pairs of points of an image, leads to a new *'virtual'* view of the object under orthographic projection. In the case of perspective projection, only D_1 is a *legal* transformation. Notice that symmetric pairs are the elementary features in this situation and points lying on the symmetry plane are degenerate cases of symmetric pairs.

Geometrically, this analysis means that, for bilaterally symmetric objects, simple transformations of a given 2D view yield other 2D views that are *legal.* It is remarkable that, in order to apply these 2D transformations to a view, it is not necessary to know the spatial orientation of the object or even its 3D structure but only that it is bilaterally symmetric *and* the symmetric pairs of features. Each transformation corresponds to a proper rotation of a rigid 3D object followed by its orthographic projection on the image plane as shown for the transformation D_1 in Figure 1 top. The transformations are different from a 2D reflection at an axis in the image plane.

In the following, we demonstrate how these 'virtual' views contain additional information, which can be used in object recognition. Let us first point out the difference between recognizing an object and computing its Euclidean 3D structure, which is much harder. We say that an object is *recognized* if its 2D view is an element of the linear vector space $V^{2N}_{ob_i}$, the space of all possible views of a known model object, obtained by 3D linear affine transformations of the 3D model followed by orthographic projection. This definition is equivalent to the affine structure used in Koenderink and van Doorn, 1991.

In the case of orthographic projection, we will show that for a bilateral symmetric object a single 2D view (and its 'virtual' view) are sufficient for recognition. Ullman and

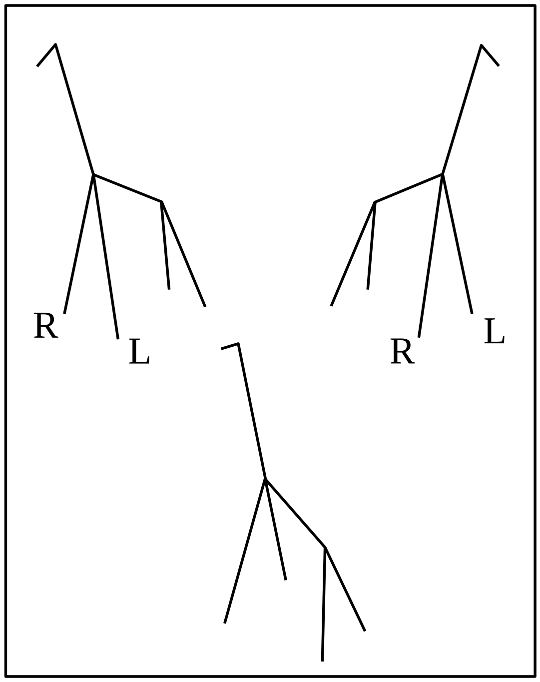

Figure 1: *Given a single 2D view (upper left), a new view (upper right) is generated under the assumption of bilateral symmetry. The two views are sufficient to verify that a novel view (second row) corresponds to the same object as the first. (The labels are only used to visualize the applied transformation D_1 between upper left and upper right, for definition see text.)*

Basri (1991) showed that the linear vector space $V_{ob_i}^{2N}$ of all possible 2D views of an object with N feature points has 6 dimensions. Their proof is equivalent to saying that $V_{ob_i}^{2N}$ consists of vectors with x and y coordinates and can be written as the direct sum $V_{ob_i}^{2N} = V_x^N \oplus V_y^N$, where V_x^N and V_y^N are non-intersecting linear subspaces, each isomorphic to $\Re^3$. This implies that the x coordinates of a 2D view are the linear combination of the x coordinates of 3 2D views and the y coordinates are the linear combination of the y coordinates of 3 2D views, the two combinations, however, being independent of each other. The 1.5 *views theorem* (Poggio 1990) proves that $V_x^N = V_y^N$ and therefore any 4 vectors from V_x^N and V_y^N are linearly dependent. So in general 2 views are sufficient to span $V_{ob_i}^{2N}$ by taking two vectors of the x and y coordinates of a first view and a third vector of the x or y coordinates of a second view. Consider now the 2D view of a bilateral symmetric object, which we assume consists of at least 4 non colinear feature points. It is easily seen that in general the vectors formed by $p_0 = (x_R, y_R, x_L, y_L)$ and $p_1 = (-x_L, y_L, -x_R, y_R)$ are linearly independent. Only for "accidental" views, like the perfectly "frontal" view or the perfect "side" view, the vectors of the x and y components are linearly dependent.

2.1.2 A Recognition Algorithm

A single 2D model view together with the knowledge that the object is bilaterally symmetric can be used for recognition in the following way.

1. Take $\mathbf{x}_1$ and $\mathbf{y}_1$ (the vectors of the x and y coordinates of the n feature points) from the available view and generate a third vector $\mathbf{x}_2$ (or $\mathbf{y}_2$) by applying the symmetry transformation D_1 to all pairs of symmetric points.

2. Make a $2n \times 6$ matrix B with its 6 columns representing a basis for $V_{obj}^{2N} = V_x^N \oplus V_y^N$. An explicit form of B is

$$B = \begin{pmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \mathbf{y}_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathbf{x}_1 & \mathbf{x}_2 & \mathbf{y}_1 \end{pmatrix}$$

3. Check that B is full rank (for instance $(B^T B)^{-1}$ exists). This is equivalent to testing for "accidental" views.

4. A novel view $\mathbf{t}$ (we assume here that the first n components are the x coordinates followed by n y coordinates) of the same object must be in the space spanned by the columns of B, and therefore must satisfy

$$\mathbf{t} = B\alpha$$

which implies (since $(B^T B)^{-1}$ exists)

$$\mathbf{t} = B(B^T B)^{-1} B^T \mathbf{t} \tag{2}$$

B can then be used to check whether $\mathbf{t}$ is a view of the correct object or not, by checking if $\|\mathbf{t} - B(B^T B)^{-1} B^T \mathbf{t}\| = 0$ (an additional test for rigidity may also be applied, if desired, to the three available views). Figure 1 shows the results of using this technique to recognize simple pipe-cleaner animals.

Notice that bilateral symmetry provides from one 2D view a total of four 2D views (image plane rotations not included), each corresponding to a different rotation of the original 3D object. Two of the four views are linearly independent (two linearly-independent vectors of the x coordinates and two for the y coordinates). The results of Shashua (1993) in combination with the virtual views prove even in the perspective case the existence of an projective invariant for bilateral symmetric objects. For recognition functions and projective invariants of symmetric objects see also Moses and Ullman (1991) and Rothwell *et al.* (1993).

2.2 Correspondence and Bilateral Symmetry

Let us suppose that the correspondence of four non coplanar points (or more) between a model view (with the pairs of symmetric feature-points already identified) and a novel view is given. Then the following epipolar line argument can be applied separately to each of the two views generated by the model view under the assumption of bilateral symmetry (see equation 1). The x, y coordinates of corresponding points in two images of an object undergoing an affine transformation are linearly dependent, that is

$$\alpha_1 \mathbf{x}_1 + \beta_1 \mathbf{y}_1 + \alpha_2 \mathbf{x}_2 + \beta_2 \mathbf{y}_2 = 0.$$

For each point (x_1, y_1) in the model view the corresponding point (x, y) in the novel view then satisfies the two equations:

$$y = mx + A \qquad \textit{and} \qquad y = m'x + A'$$

with $m = -\alpha_2/\beta_2$ and $A = -(\alpha_1 x_1 + \beta_1 y_1)/\beta_2$ and m', A' determined by the virtual view. Therefore (x, y) is uniquely determined (apart special cases) as

$$y = \frac{m'A - mA'}{m' - m} \quad , \quad x = \frac{A' - A}{m - m'} \quad .$$

Thus, *the correspondence of four non-coplanar points between two 2D views of a bilateral symmetric object (undergoing a uniform affine transformation) uniquely determines the correspondence of all other points.*

In the case of occlusions, correspondence of singular points in the novel view can be established as long as the pair of symmetric points is identified in the model view. When full correspondence between the model and novel view is established, any structure from motion algorithm can be used to compute the Euclidean structure related to the two views and the assumption of symmetry.

2.3 Euclidean Structure from One 2D Model View

Suppose, as before, that we have a single 2D view of an object. Assume further that we hypothesize (correctly) that the object is two-fold bilaterally symmetric and that symmetric quadruples can be identified. These are the "elementary" features in this situation, since any point, that is not on both symmetry planes corresponds to three other points. We define an object to be two-fold bilaterally symmetric if the following transformations of any 2D view of a feature quadruple yield *legal views* of the quadruple, that is orthographic projections of rigid rotations of the object:

$$D_{21}\mathbf{q}_{quadr} = \mathbf{q}^1_{quadr} \tag{3}$$

$$D_{22}\mathbf{q}_{quadr} = \mathbf{q}^2_{quadr} \tag{4}$$

with

$$\mathbf{q}_{quadr} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} \quad , \quad \mathbf{q}^1_{quadr} = \begin{pmatrix} -x_2 \\ -x_1 \\ -x_4 \\ -x_3 \\ y_2 \\ y_1 \\ y_4 \\ y_3 \end{pmatrix} \quad \textit{and} \quad \mathbf{q}^2_{quadr} = \begin{pmatrix} x_4 \\ x_3 \\ x_2 \\ x_1 \\ -y_4 \\ -y_3 \\ -y_2 \\ -y_1 \end{pmatrix} .$$

These 3 views are independent apart from special views, such as accidental views (see previous section). Thus the above definition of symmetry provides a way to generate two

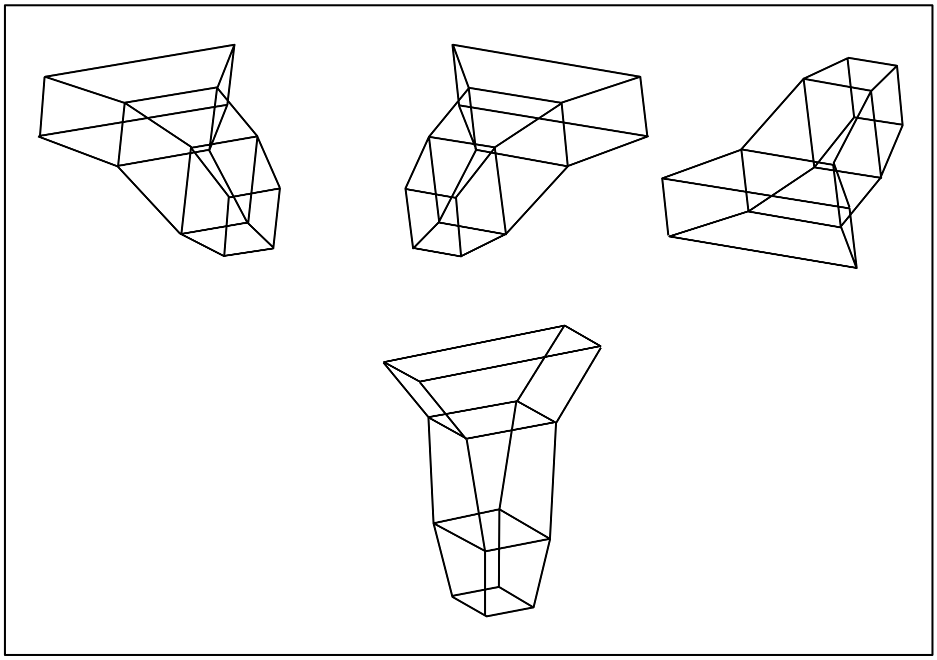

Figure 2: *A single 2D view (upper left) of a two-fold bilateral symmetric object can generate additional views (upper center and right) using the symmetric properties of the object. Those three views are sufficient to compute the 3D structure, as indicated in the second row where we show a new projection of the 3D structure computed from the three views above.*

additional views from the given single view, unless $\mathbf{q}_{quadr}$ is a view which is left invariant by at least one of the symmetry transformations. This is the case, for instance, for exactly frontal views. The same argument can be repeated for all symmetric quadruples. These transformations are the same transformations from the previous section applied to both symmetries.

Thus, these transformations yield in the generic case 3 independent views of the object (the symmetry yields a total of 16 views, representing 16 different orientations of the object, which span the 6-dimensional viewing space of the object). One can verify that standard structure-from-motion techniques (Huang and Lee, 1989; see also Ullman, 1979) can be applied to conclude that structure is uniquely determined, except for a reflection about the image plane. The matrix defined by Weinshall (1993) to compute an object invariant is full rank in this case; it is, however, rank deficient for simple bilateral symmetry. Using a different approach, the pairwise comparison technique of Kontsevich (1993) comes to a similar result. The following holds:

Given a single 2D orthographic view of a two-fold bilateral symmetric object (with at least two symmetric, nondegenerate quadruple features containing a total of at least four non-coplanar points) the corresponding structure is uniquely determined up to a reflection about the image plane.

In addition, the following results can be easily derived:

1. 3D structure can be obtained from two 2D views of a bilateral symmetric object (see figure 2).

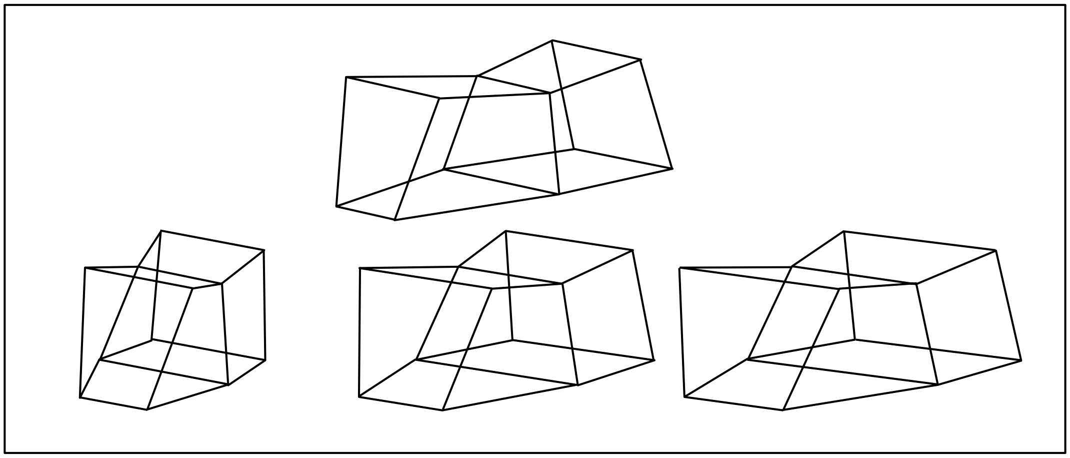

Figure 3: *A single 2D view (upper row) of a bilateral symmetric object can be generated by different bilateral symmetric 3D objects. The three objects projected in the second row all generate the 2D view of the first row after a rotation of* 20° *around the vertical axis.*

2. Structure cannot be uniquely obtained from a single 2D view of a bilateral symmetric object. So a single 2D view of a bilateral symmetric object can be generated by different bilateral symmetric objects (see figure 3).

3 Discussion

Exploiting knowledge about the symmetry of an object, recognition is possible from a single view. The geometric constraints of symmetry allow the generation of additional legal views from the given one, using 2D image transformations. This can be done for all non occluded pairs of symmetric points, without knowing the orientation of the symmetry plane or the camera position. For bilaterally symmetric objects we proved a single view is sufficient for recognition from a different view. For two-fold bilateral symmetric objects the 3D structure can be computed from a single view. Here are some implications of our results:

- *Exact frontal model views should be avoided*

  The results about bilateral symmetry imply that one should avoid using a model view which is a fixed point of the symmetry transformations (since the transformation of it generates an identical new view). In the case of faces, this implies that the model view in the data base should not be an exactly frontal view. Psychophysical evidence supporting this point is given by Schyns and Bülthoff, (1994).

- *A symmetry of order higher than bilateral allows recovery of structure from one 2D view*

  Our results imply that even in absence of 3D cues (such as shading, perspective, texture etc.), an object symmetry of sufficiently high order may provide structure from a single view. An interesting conjecture is that human perception may be biased to impose a symmetry assumption (in the absence of other evidence to the contrary), in order to compute structure.

- *A new algorithm for computing structure from single views of polyedric objects*

  For line drawings, Marrill (1991), Sinha and Adelson (1993) proposed an iterative algorithm that is capable of recovering structure from single views. Our result on structure-from-1-view may explain some of these results in terms of the underlying algebraic structure induced by symmetry properties. It also yields a new non-iterative algorithm for the recovery of structure since it provides (once symmetric n-tuples are identified) a simple algorithm generating a total of three linearly-independent views to which any of the classical Structure-from-Motion algorithms can be applied, including the recent linear algorithms (Huang and Lee, 1989). It remains an open question to characterize the connection between the minimization principle of Marrill-Sinha and our internal structure constraints. Especially in the case of bilateral symmetric objects their principle might help to understand, which constraints are used from human observers to disambiguate views as shown in figure 3.

- *"Virtual" views and image based object recognition*

  For image-based recognition systems (Poggio and Edelman, 1990), the possibility of generating additional views for objects minimizes the number of necessary example views. In the case of symmetric objects, the image transformations related to rotations in 3D space can be derived directly from the 3D structure of the class of symmetric objects. In the case of nonlinear transformations, the related image transformations have to be approximated from prototypical views of a class of objects. The approach by Beymer *et al.*(1993) demonstrates how this can be done for pose and expressions of human faces. Novel grey-level images, related to changes of facial expression, can be generated from a single image when applying the appropriate image transformation.

- *Psychophysical results on object recognition*

  It is intriguing to speculate about relations between the known human abilities of detecting symmetries and the human tendencies of hypothesizing symmetry in visual perception. There is evidence on spontaneous generalization to left-right reversal in humans and even simpler visual systems (see Rock *et al.*, 1989; Sutherland, 1960; Young, 1964). Our theory offers a simple explanation of these effects as a by-product of a mechanism optimized for the recognition of 3D objects. Thus, visual recognition of 3D objects may be the main reason for the well known sensitivity of visual systems to bilateral symmetry of 3D objects and 2D patterns.

  The results found in psychophysical experiments on object recognition (Vetter *et al.*, 1994) are consistent with our theoretical predictions for symmetric objects. Based on a single training view, the generalization performance for novel views is significantly better for symmetric objects than for non-symmetric objects. In contrast to the non-symmetric objects, the generalization field of symmetric objects showed additional peaks of good recognition performance. These additional peaks were in all cases at the location of the virtual views. It is not yet clear in what way the visual system uses symmetry: instead of creating "explicit" virtual views the system may discover and use symmetry-based features that are view invariant.

Several open questions remain. How does a visual system, natural or artificial, detect symmetric pairs of features of a 3D object, a task which is in general quite different from symmetry detection in a 2D pattern. What are the optimal cues leading to the assumption of symmetry, since it is not possible to prove the symmetry when only a single view is given? In some cases (e.g., line drawings of geometric objects), algorithms capable of identifying feature points likely to be symmetric should be feasible, since all pairs of symmetric points in one view obey to the same epipolar line constraint. In other cases additional information may be available (e.g. once the two eyes are identified as eyes, it is known that they represent a symmetric pair). There the knowledge about the symmetry of the object class can help to establish the correspondence between symmetric feature pairs. Another question which is open is how to extend our approach of using 2D image transformations to geometric constraints other than bilateral symmetry.

Acknowledgments
We are grateful to H.H. Bülthoff, F. Girosi, P. Sinha and S. Ullman for useful discussions and suggestions. This research is sponsored by grants from the Office of Naval Research under contracts N00014-91-J-1270 and N00014-92-J-1879; by a grant from the National Science Foundation under contract ASC-9217041 (including funds from DARPA provided under the HPCC program).

Tomaso Poggio is supported by the Uncas and Helen Whitaker Chair at the Whitaker College, Massachusetts Institute of Technology. Thomas Vetter was supported by a post-doctoral fellowship from the Deutsche Forschungsgemeinschaft (Ve 135/1-1).

References

Beymer, D., Shashua, A., and Poggio, T. (1993). Example-based image anaysis and synthesis. A.I. Memo No. 1431, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Edelman, S. and Poggio, T. (1992). Bringing the Grandmother back into the picture: a memory-based view of object recognition. *Int. J. Pattern Recog. Artif. Intell.*, 6.

Faugeras, O. (1992). What can be seen in three dimensions with an uncalibrated stereo rig? *Proceedings ECCV*, pages 563–578.

Hartley, R. I., Gupta, R., and Chang, T. (1992). Stereo from uncalibrated cameras. *Proceedings CVPR92*, pages 761–764.

Huang, T. and Lee, C. (1989). Motion and structure from orthographic projections. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2(5):536–540.

Koenderink, J. and van Doorn, A. (1991). Affine stucture from motiom. *Journal of the Optical Society of America*, 8(2):377–385.

Kontsevich, L. L. (1993). Pairwise comparison technique: a simple solution for deph reconstruction. *J.Opt.Soc.Am.A*, 10(6):1129–1135.

Marill, T. (1991). Emulating the human interpretation of line-drawings as three-dimensional objects. *International Journal of Comuter Vision*, 6:147–161.

Moses, Y. and Ullman, S. (1991). Limitations of non model-based recognition schemes. A.I.Memo 1301, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Poggio, T. (1990). 3D object recognition: on a result by Basri and Ullman. Technical Report # 9005–03, IRST, Povo, Italy.

Poggio, T. and Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, 343:263–266.

Poggio, T. and Vetter, T. (1992). Recognition and structure from one 2D model view: observations on prototypes, object classes, and symmetries. A.I. Memo No. 1347, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Rock, I., Wheeler, D., and Tudor, L. (1989). Can we imagine how objects look from other viewpoints? *Cognitive Psychology*, 21:185–210.

Rothwell, C., Forsyth, D., Zissermann, A., and Mundy, J. (1993). Extracting projective structure from single perspective views of 3D point sets. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 573–582, Berlin, Germany.

Schyns, P. and Bülthoff, H. (1993). Sparse observations on cortical mechanisms for object recognition and learning. A.I. Memo No. 1404, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Shashua, A. (1993). Projective depth: A geometric invariant for 3D reconstruction. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 583–590, Berlin, Germany.

Sinha, P. and Adelson, E. H. (1993). Recovering 3D shapes from 2D line-drawings. In *Intelligent Robotics: Procs. of the IEEE Int.Symp. on Intelli. Robotics*, Bangalore, India.

Sutherland, N. (1960). Visual discrimination of orientation by *octopus*: mirror images. *Brit. J. Psychol.*, 51:9–18.

Ullman, S. (1979). *The interpretation of visual motion.* MIT Press, Cambridge, MA.

Ullman, S. and Basri, R. (1991). Recognition by linear combinations of models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:992–1006.

Vetter, T., Poggio, T., and Bülthoff, H. (1994). The importance of symmetry and virtual views in three-dimensional object recognition. *Current Biology*, 4:18–23.

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

Young, J. (1964). *A model of the brain.* Oxford University Press, Oxford.