

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY
and
CENTER FOR BIOLOGICAL AND COMPUTATIONAL LEARNING
WHITAKER COLLEGE

A.I. Memo No. 1448
C.B.C.L. Paper No. 85

June, 1994

The Quadric Reference Surface: Theory and Applications

Amnon Shashua and Sebastian Toelg

Abstract

The conceptual component of this work is about “reference surfaces” which are the dual of reference frames often used for shape representation purposes. The theoretical component of this work involves the question of whether one can find a unique (and simple) mapping that aligns two arbitrary perspective views of an opaque textured quadric surface in 3D, given (i) few corresponding points in the two views, or (ii) the outline conic of the surface in one view (only) and few corresponding points in the two views. The practical component of this work is concerned with applying the theoretical results as tools for the task of achieving full correspondence between views of arbitrary objects.

Short version of this manuscript appears in the Proceedings of ECCV'94, Stockholm, Sweden.

Copyright © Massachusetts Institute of Technology, 1994

This report describes research done within the Center for Biological and Computational Learning in the Department of Brain and Cognitive Sciences and at the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. This research is sponsored by grants from the Office of Naval Research under contracts N00014-92-J-1879 and N00014-93-1-0385; and by a grant from the National Science Foundation under contract ASC-9217041 (this award includes funds from ARPA provided under the HPCC program). Additional support is provided by the North Atlantic Treaty Organization, ATR Audio and Visual Perception Research Laboratories, Mitsubishi Electric Corporation, Sumitomo Metal Industries, and Siemens AG. Support for the A.I. Laboratory's artificial intelligence research is provided by ARPA contract N00014-91-J-4038. A. Shashua is supported by a McDonnell-Pew postdoctoral fellowship from the department of Brain and Cognitive Sciences. S. Toelg was supported by a postdoctoral fellowship from the Deutsche Forschungsgemeinschaft. S. Toelg is currently at the Center for Automation Research at the University of Maryland at College Park.

1 Introduction

This paper has three main goals, which in a way stand on their own. First, to support and extend the concept of a “reference surface”, which currently exists only in rudimentary form. Second, to introduce a natural application for more advanced reference surfaces for the purpose of achieving visual correspondence or registration across images (views) of 3D scenes. Third, to introduce new theoretical results on a specific class of reference surfaces, the quadrics. The crux of this work is embedded in the third goal, yet we emphasize briefly the first two goals as they provide a context for one particular use of our results, and reasons for pushing further along these and similar lines (see also the discussion in Section 6).

Reference surfaces are simply the dual of reference frames used for shape representation. An object in 3D space is represented relative to some frame. For example, if we model an object as a collection of points, then in affine space a minimal frame must consist of four points in general position; in projective space a minimal frame consists of five points in general position. In a dual manner, in affine space a reference plane is minimally necessary for shape representation; in projective space we have the tetrahedron of reference. Work along the lines of representing shape using minimal frame configurations and recovery from views can be found in [9, 8, 4, 27, 28], and in further references therein.

As long as we use the minimal configuration of points for representing shape, there is no practical reason to distinguish between reference frames and reference surfaces. The distinction becomes useful, as we shall see later, when we choose non-minimal frames; their dual corresponds to non-planar reference surfaces. Before we elaborate further on the duality between reference frames and reference surfaces, it would be useful to consider a specific application in which the notion of reference surfaces appears explicitly.

Consider the problem of achieving correspondence, or optical flow as it is known in the motion literature. The task is to recover the 2D displacement vector field between points across two images, in particular in the case where the two images are two distinct views of some 3D object or scene. Typical applications for which often full correspondence (that is correspondence for all image points) is initially required include the measurement of motion, stereopsis, structure from motion, 3D reconstruction from point correspondences, and more recently, visual recognition, active vision and computer graphics animation.

The concept of reference surfaces becomes relevant in this context when we consider the correspondence problem as composed of two stages: (i) a nominal transformation phase due to a reference surface, and (ii) recovery of a residual field (cf. [24, 25, 2]). In other words, we envision some virtual reference surface on the object and project the object onto that surface along the lines of sight associated with one of the views. As a result, we have two surfaces, the original object and a virtual object (the reference surface). The correspondence field between the two views generated by the virtual surface can be characterized by a closed-form transforma-

tion (the “nominal transformation”). The differences between the corresponding points coming from the original surface and the corresponding points coming from the virtual surface are along epipolar lines and are small in regions where the reference surface lies close to the original surface. These remaining displacements are referred to as the “residual displacements”. The residuals are recovered using instantaneous spatio-temporal derivatives of image intensity values along the epipolar lines (see Fig. 1).

As a simple example, consider the case where the reference surface is a plane. It is worth noting that planar reference surfaces are also found in the context of navigation and obstacle detection in active vision applications [29, 30, 13] as well as in infinitesimal motion models for visual reconstruction [7, 10, 23]. A planar reference surface corresponds to the dual case of shape representation under parallel projection (cf. [9]), or relative affine structure under perspective projection [26, 28]. In other words, a nominal transformation is either a 2D affine transformation or a 2D projective transformation, depending on whether we assume an orthographic or perspective model of projection. The magnitude of the residual field is thus small in image regions that correspond to object points that are close to the reference plane, and the magnitude is large in regions that correspond to object points that are far away from the reference plane. This is demonstrated in Fig. 2. The top row displays show two views of a face obtained by rotation of the head approximately around the vertical axis of the neck. Three points were chosen (two eyes and the right mouth corner) for computing the nominal transformation. The overlay of the second view and the transformed first view demonstrate (bottom row) that the central region of the face is brought closer at the expense of regions near the boundary, which correspond to object points that are far away from the virtual plane passing through both eyes and the mouth corner.

This example naturally suggests that a nominal transformation based on placing a virtual quadric reference surface on the object would give rise to a smaller residual field — for this particular class of objects. A quadric reference surface is a natural extension of the planar case and, as the example above demonstrates, may be a useful tool for the application of visual correspondence.

In terms of duality between frames for shape representation and reference surfaces, the quadric reference frame will require a non-minimal configuration (of points and other forms). This configuration can also serve as a frame for shape representation, but the property we emphasize here is the use of its dual — the quadric reference surface.

The theoretical component of this work is therefore concerned with establishing a quadric reference surface from image information across two views. We start by addressing the following questions: First, given any two views of some unknown textured opaque quadric surface in 3D projective space \mathcal{P}^3 , is there a finite number of corresponding points across the two views that uniquely determine all other correspondences coming from points on the quadric? Second, can the unique mapping be de-




Figure 1: Schematic illustration of the main concepts. The object is projected onto a virtual quadric along the line of sight. Points on the quadric are then projected onto the second image plane. The deviation of the object from the virtual quadric is a measure of shape (the quadric is a reference surface); the transformation $T(p)$ due to the quadric is the “nominal quadratic transformation”; the displacement between p' and $T(p)$ is along epipolar lines and is called “residual displacement”. This paper is about deriving (Theorems 1, 4) general methods for recovering $T(p)$ given a small number of corresponding points across the two views (four points and the epipoles are sufficient).

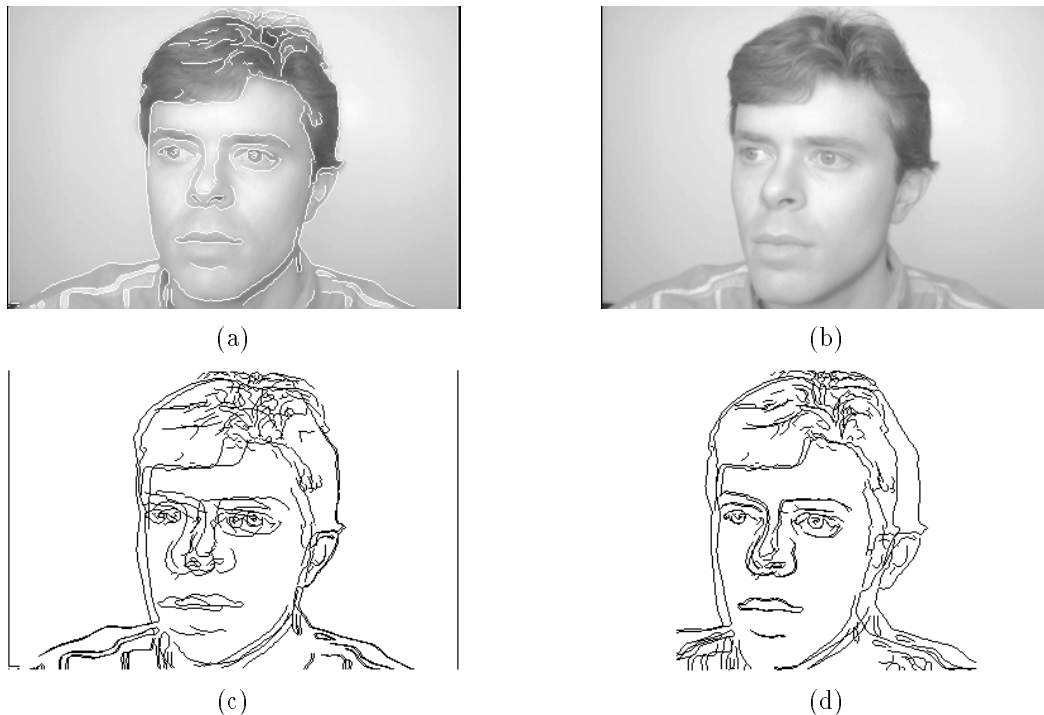


Figure 2: The case of a planar reference surface. (a), (b) are images of two views of a face, first view ψ_1 on the left and second view ψ_2 on the right. Edges are superimposed on (a) for illustrative purposes. (c) overlaid edges of ψ_1 and ψ_2 . (d) The residual displacements (see text and Fig. 1) resulting from a planar reference surface. The planar nominal transformation is the 2D affine transformation determined by three corresponding points across the two eyes and the right mouth corner of the face. Notice that the displacements across the center region of the face are reduced, at the expense of the peripheral regions which are taken farther apart.

terminated by the outline conic in one of the views (projection of the rim) and a smaller number of corresponding points? A constructive answer to these questions readily suggests that we can associate a virtual quadric surface with any 3D object (not necessarily itself a quadric) and use it for describing shape, but more importantly, for achieving full correspondence between the two views.

On the conceptual level we propose combining geometric constraints, captured from knowledge of a small number of corresponding points (manually given, for example), and photometric constraints captured by the instantaneous spatio-temporal changes in image brightness (conventional optical flow). The geometric constraints we propose are related to the virtual quadric surface mentioned above. These constraints lead to a transformation (a nominal quadratic transformation) that is applied to one of the views with the result of bringing both views closer together. The remaining displacements (residuals) are recovered either by optical flow techniques using the spatial and temporal derivatives of image brightness or by correlation of image patches.

2 Notation

We consider object space to be the 3D projective space \mathcal{P}^3 , and image space to be the 2D projective space \mathcal{P}^2 — both over the field \mathcal{C} of complex numbers. Views are denoted by ψ_i , indexed by i . The epipoles are denoted by $v \in \psi_1$ and $v' \in \psi_2$, and we assume their locations are known (for methods, see [4, 5, 27, 28, 12], for example, and briefly later in the text). The symbol \cong denotes equality up to scale, GL_n stands for the group of $n \times n$ matrices, PGL_n is the group defined up to scale, and $SPGL_n$ is the symmetric specialization of PGL_n .

3 The Quadric Reference Surface I: Points

We start with recovering the parameters of a quadric, modeled as a cloud of points, from two of its projections. The problem is straightforward if the two projection centers are on the surface (Result 1). The general case (Theorem 1) is also made easier by resorting to projective reconstruction via a simple and convenient parameterization of space (Lemma 1). Guaranteeing uniqueness of the mapping between the two views of the quadric is somewhat challenging because the ray from a projection center generally intersects the quadric at two points. This situation is disambiguated by combining an “opacity” assumption (Definition 1) with the parameterization used for recovering the quadric parameters (Lemma 2). Finally, as a byproduct of these derivations, one can readily obtain quantitative and simple measures relating the projection centers and the family of quadrics passing through configurations of eight points (Theorem 2). This may have applications in the analysis of “critical surfaces” (Corollary 1).

Result 1 *Given two arbitrary views $\psi_1, \psi_2 \subset \mathcal{P}^2$ of a quadric surface $Q \in \mathcal{P}^3$ with centers of projection at $O, O' \in \mathcal{P}^3$, and $O, O' \in Q$, then five corresponding points across the two views uniquely determine all other correspondences.*

Proof. Let (x_0, x_1, x_2) and (x'_0, x'_1, x'_2) be coordinates of ψ_1 and ψ_2 , respectively, and (z_0, \dots, z_3) be coordinates of Q . Let $O = (0, 0, 0, 1)$; then the quadric surface may be given as the locus $z_0z_3 - z_1z_2 = 0$, and ψ_1 as the projection from $O = (0, 0, 0, 1)$ onto the plane $z_3 = 0$. In case where the centers of projection are on Q , the line through O meets Q in exactly one other point, and thus the mapping $\psi_1 \mapsto Q$ is generically one-to-one, and so has a rational inverse: $(x_0, x_1, x_2) \mapsto (x_0^2, x_0x_1, x_0x_2, x_1x_2)$. Because all quadric surfaces of the same rank are projectively equivalent, we can perform a similar blow-up from ψ_2 with the result (x'_0, x'_1, x'_2, x'_3) . The projective transformation $A \in PGL_4$ between the two representations of Q can then be recovered from five corresponding points between the two images. \square

This result does not hold when the centers of projection are not on the quadric surface. This is because the mapping between Q and \mathcal{P}^2 is not one-to-one (a ray through the center of projection meets Q in two points), and therefore, a rational inverse does not exist. We are interested in establishing a more general result that applies when the centers of projection are not on the quadric surface. One way to enforce a one-to-one mapping is by making “opacity” assumptions, defined below.

Definition 1 (Opacity Constraint) *Given an object $Q = \{P_1, \dots, P_n\}$, we assume there exists a plane through the camera center O that does not intersect any of the chords P_iP_j (i.e., Q is observed from only one “side” of the camera). Furthermore, we assume that the surface is opaque, which means that among all the surface points along a ray from O , the closest point to O is the point that also projects to the second view (ψ_2). The first constraint, therefore, is a camera opacity assumption, and the second constraint is a surface opacity assumption — which together we call the opacity constraint.*

With an appropriate parameterization of \mathcal{P}^3 we can obtain the following result:

Theorem 1 *Given two arbitrary views $\psi_1, \psi_2 \subset \mathcal{P}^2$ of an opaque quadric surface $Q \in \mathcal{P}^3$; then nine corresponding points across the two views uniquely determine all other correspondences.*

The following auxiliary propositions are used as part of the proof.

Lemma 1 (Relative Affine Parameterization)

Let p_o, p_1, p_2, p_3 and p'_o, p'_1, p'_2, p'_3 be four corresponding points coming from four non-coplanar points in space. Let A be a collineation (homography) of \mathcal{P}^2 determined by the equations $Ap_j \cong p'_j$, $j = 1, 2, 3$, and $Av \cong v'$. Finally let v' be scaled such that $p'_o \cong Ap_o + v'$. Then, for any point $P \in \mathcal{P}^3$ projecting onto p and p' , we have

$$p' \cong Ap + kv'. \quad (1)$$

The coefficient $k = k(p)$ is independent of ψ_2 , i.e., is invariant to the choice of the second view, and the projective coordinates of P are $(x, y, 1, k)^T$.

The lemma, its proof and its theoretical and practical implications are discussed in detail in [26, 28]. The scalar k is called a *relative affine invariant* and can be computed with the assistance of a second arbitrary view ψ_2 . In a nutshell, a representation \mathcal{R}_o of \mathcal{P}^3 is chosen such that the projection center O of the first camera position is part of the reference frame (of five points). The matrix A is the 2D projective transformation due to the plane π passing through the object points P_1, P_2, P_3 , i.e., for any $P \in \pi$ we have $p' \cong Ap$. The representation \mathcal{R}_o is associated with $[x, y, 1, k]$ where k vanishes for all points coplanar with π , which means that π is the plane at infinity under the representation \mathcal{R}_o . Finally, the transformation between \mathcal{R}_o and the representation \mathcal{R} as seen from any other camera position (uncalibrated), can be described by an element of the affine group, i.e., the scalar k is an affine invariant relative to \mathcal{R}_o .

Proof of Theorem: From Lemma 1, any point P can be represented by the coordinates $(x, y, 1, k)$ and k can be computed from Equation 1. Since Q is a quadric surface, there exists $H \in SPGL_4$ such that $P^T H P = 0$, for all points P of the quadric. Because H is symmetric and determined up to scale, it contains only nine independent parameters. Therefore, given nine corresponding image points we can solve for H as a solution of a linear system; each corresponding pair p, p' provides one linear equation in H of the form $(x, y, 1, k)H(x, y, 1, k)^T = 0$.

Given that we have solved for H , the mapping $\psi_1 \mapsto \psi_2$ due to the quadric Q can be determined uniquely (i.e., for every $p \in \psi_1$ we can find the corresponding $p' \in \psi_2$) as follows. The equation $P^T H P = 0$ gives rise to a second order equation in k of the form $ak^2 + b(p)k + c(p) = 0$, where the coefficient a is constant (depends only on H) and the coefficients b, c depend also on the location of p . Therefore, we have two solutions for k , and by Equation 1, two solutions for p' . The two solutions for k are $k^1, k^2 = \frac{-b \pm r}{2a}$, where $r = \sqrt{b^2 - 4ac}$. The finding, shown in the next auxiliary lemma, is that if the surface Q is opaque, then the sign of r is fixed for all $p \in \psi_1$. Therefore, the sign of r for p_o that leads to a positive root (recall that $k_o = 1$) determines the sign of r for all other $p \in \psi_1$. \square

Lemma 2 *Given the opacity constraint, the sign of the term $r = \sqrt{b^2 - 4ac}$ is fixed for all points $p \in \psi_1$.*

Proof. Let P be a point on the quadric projecting onto p in the first image, and let the ray \overline{OP} intersect the quadric at points P^1, P^2 , and let k^1, k^2 be the roots of the quadratic equation $ak^2 + b(p)k + c(p) = 0$. The opacity assumption is that the intersection closer to O is the point projecting onto p and p' .

Recall that P_o is a point (on the quadric in this case) used for setting the scale of v' (in Equation 1), i.e., $k_o = 1$. Therefore, all points that are on the same side of π as P_o have positive k associated with them, and vice versa (similar logic was used in [21] for convex-hull computations). There are two cases to be considered: either P_o is between O and π (i.e., $O < P_o < \pi$), or π is between O and P_o (i.e., $O < \pi < P_o$) — that is O and P_o are on opposite sides of π . In the first case,

if $k^1 k^2 \leq 0$ then the non-negative root is closer to O , i.e., $k = \max(k^1, k^2)$. If both roots are negative, the one closer to zero is closer to O , again $k = \max(k^1, k^2)$. Finally, if both roots are positive, then the larger root is closer to O . Similarly, in the second case we have $k = \min(k^1, k^2)$ for all combinations. Because P_o can satisfy either of these two cases, the opacity assumption then gives rise to a consistency requirement in picking the right root: either the maximum root or the minimum root should be uniformly chosen for all points. \square

In Section 5 we will show that Theorem 1 can be used to surround an arbitrary 3D surface by a virtual quadric, i.e., to create quadric reference surfaces, which in turn can be used to facilitate the correspondence problem between two views of a general object. The remainder of this section takes Theorem 1 further to quantify certain useful relationships between the centers of two cameras and the family of quadrics that pass through arbitrary configurations of eight points whose projections on the two views are known.

Theorem 2 *Given a quadric surface $Q \subset \mathcal{P}^3$ projected onto views $\psi_1, \psi_2 \subset \mathcal{P}^2$, with centers of projection $O, O' \in \mathcal{P}^3$, there exists a parameterization of the image planes ψ_1, ψ_2 that yields a representation $H \in SPGL_4$ of Q such that $h_{44} = 0$ when $O \in Q$, and the sum of the elements of H vanishes when $O' \in Q$.*

Proof. The re-parameterization described here was originally introduced in [26] as part of the proof of Lemma 1. We first assign the standard coordinates in \mathcal{P}^3 to three points on Q and to the two camera centers O and O' as follows. We assign the coordinates $(1, 0, 0, 0)$, $(0, 1, 0, 0)$, $(0, 0, 1, 0)$ to P_1, P_2, P_3 , respectively, and the coordinates $(0, 0, 0, 1)$, $(1, 1, 1, 1)$ to O, O' , respectively. By construction, the point of intersection of the line $\overline{OO'}$ with π has the coordinates $(1, 1, 1, 0)$.

Let P be some point on Q projecting onto p, p' . The line \overline{OP} intersects π at the point $(\alpha, \beta, \gamma, 0)$. The coordinates α, β, γ can be recovered (up to scale) by the mapping $\psi_1 \mapsto \pi$, as follows. Given the epipoles v and v' , we have by our choice of coordinates that p_1, p_2, p_3 and v are projectively (in \mathcal{P}^2) mapped onto $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$, $e_3 = (0, 0, 1)$ and $e = (1, 1, 1)$, respectively. Therefore, there exists a unique element $A_1 \in PGL_3$ that satisfies $A_1 p_j \cong e_j$, $j = 1, 2, 3$, and $A_1 v \cong e$. Let $A_1 p = (\alpha, \beta, \gamma)$. Similarly, the line $\overline{O'P}$ intersects π at $(\alpha', \beta', \gamma', 0)$. Let $A_2 \in PGL_3$ be defined by $A_2 p'_j \cong e_j$, $j = 1, 2, 3$, and $A_2 v' \cong e$. Let $A_2 p' = (\alpha', \beta', \gamma')$.

It is easy to see that $A \cong A_2^{-1} A_1$, where A is the collineation defined in Lemma 1. Likewise, the homogeneous coordinates of P are transformed into $(\alpha, \beta, \gamma, k)$. With this new coordinate representation the assumption $O \in Q$ translates to the constraint that $h_{44} = 0$ ($((0, 0, 0, 1)H(0, 0, 0, 1)^T = 0)$), and the assumption $O' \in Q$ translates to the constraint $(1, 1, 1, 1)H(1, 1, 1, 1)^T = 0$. Note also that $h_{11} = h_{22} = h_{33} = 0$ due to the assignment of standard coordinates to P_1, P_2, P_3 . \square

Corollary 1 *Theorem 2 provides a quantitative measure of the proximity of a set of eight 3D points, projecting onto two views, to a quadric that contains both centers of projection.*

Proof. Given eight corresponding points we can solve for H with the constraint $(1, 1, 1, 1)H(1, 1, 1, 1)^\top = 0$. This is possible since a unique quadric exists for any set of nine points in general position (the eight points and O'). The value of h_{44} is then indicative of how close the quadric is to the other center of projection O . \square

Note that when the camera center O is on the quadric, then the leading term of $ak^2 + b(p)k + c(p) = 0$ vanishes ($a = h_{44} = 0$), and we are left with a linear function of k . We see that it is sufficient to have a bi-rational mapping between Q and only one of the views without employing the opacity constraint. This is because of the asymmetry introduced in our method: the parameters of Q are reconstructed with respect to the frame of reference \mathcal{R}_o which includes the first camera center (i.e., relative affine reconstruction in the sense of [26]) rather than reconstructed with respect to a purely object-based frame (i.e., all five reference points coming from the object). Also note the importance of obtaining quantitative measures of proximity of an eight-point configuration of 3D points to a quadric that contains both centers of projection; this is a necessary condition for observing a “critical surface”. A sufficient condition is that the quadric is a hyperboloid of one sheet [6, 15]. Theorem 2 provides, therefore, a tool for analyzing part of the question of how likely are typical imaging situations within a “critical volume”.

4 The Quadric Reference Surface II: Conic + Points

The previous section dealt with the problem of recovering a unique mapping between two views of an opaque quadric from point correspondences. Here we deal with a similar problem, but in addition to observing point correspondences, we observe the outline (the projection of the rim) of the quadric in one of the images. On the theoretical level, this case is challenging because we are not using the reconstruction paradigm as in Theorem 1, simply because we are observing the outline in one view only. In this case the opacity constraint, as manifested computationally in Lemma 2, plays a significant role at the level of recovering the quadric’s parameters; whereas in the previous section the opacity constraint was used only for disambiguating the mapping between the two views given the quadric’s parameters. On the practical level, this case provides significant advantages over the previous case of using point matches only (see later in Section 5).

Theorem 3 (Outline Conic) *Let $H \in SPGL_4$ represent a quadric surface $Q \subset \mathcal{P}^3$, and compose H as*

$$H = \begin{pmatrix} E & h \\ h^\top & h_{44} \end{pmatrix}, \quad (2)$$

where $E \in GL_3$ and symmetric. Let $p = (x, y, 1)^\top$ be a point (in standard coordinate representation) in a view $\psi_1 \subset \mathcal{P}^2$ of Q with projection center $O = (0, 0, 0, 1)^\top$, then

$$E' = hh^\top - h_{44}E$$

represents the outline conic (the projection of the rim) of Q in ψ_1 .

Proof. Let $P = (x, y, 1, k)^\top$ be the coordinates of points on Q . We then obtain

$$P^\top HP = p^\top Ep + 2h^\top kp + h_{44}k^2 = 0.$$

The outline conic is defined by the border between the real and complex conjugate roots of k . Thus, the roots of k are the solution to the equation $ak^2 + b(p)k + c(p) = 0$, where

$$\begin{aligned} a &= h_{44} \\ b(p) &= 2h^\top p \\ c(p) &= p^\top Ep. \end{aligned}$$

The condition for real roots, as is required for points coming from the quadric, is a non-negative discriminant $\Delta = b^2 - 4ac \geq 0$, or

$$\Delta = 4p^\top (hh^\top - h_{44}E)p \geq 0.$$

Let $E' = hh^\top - h_{44}E$. We see that the border between real and complex roots is a conic described by $p^\top E'p = 0$. \square

Theorem 4 (Outline conic and four corresponding points)

Given two arbitrary views $\psi_1, \psi_2 \subset \mathcal{P}^2$ of an opaque quadric surface $Q \in \mathcal{P}^3$, and the outline conic of Q in ψ_1 , then four corresponding points across the two views uniquely determine all other correspondences.

Proof. Let $E' \in SPGL_3$ be the representation of the given outline conic of Q in ψ_1 and let H be the representation (having the form (2)) of Q that we seek to recover. From Theorem 3 we have $E' = hh^\top - h_{44}E$. Note that if H is scaled by α , then E' is scaled by α^2 . Thus, given E' (with an arbitrary scale), we can hope to recover H at most up to a sign flip. What we need to show is that with four corresponding points, coming from a general configuration on Q , we can recover h and h_{44} (up to the sign flip). Let $P = (x, y, 1, k)^\top$ be a point on Q projecting onto $p = (x, y, 1)^\top$ in ψ_1 . We then have

$$h_{44}P^\top HP = (p, k)^\top \begin{pmatrix} hh^\top - E' & h_{44}h \\ h_{44}h^\top & h_{44}^2 \end{pmatrix} \begin{pmatrix} p \\ k \end{pmatrix} = 0,$$

which expands to

$$(p^\top h + h_{44}k)^2 = p^\top E'p = \frac{1}{4}\Delta$$

or

$$p^\top h + h_{44}k = \pm \frac{1}{2}\sqrt{\Delta}.$$

From Lemma 2, $\sqrt{\Delta}$ are either all positive or all negative; therefore if we are given four points with their corresponding k , we have exactly two solutions for (h, h_{44}) (as a solution of a linear system). The two solutions are (h, h_{44}) and $-(h, h_{44})$ and we can choose one of them arbitrarily — since any H representing Q is only determined up to scale. From Lemma 1, we can set $k = 0$ for three of the four points, and $k = 1$ to the fourth point. Finally, after recovering H , the correspondence p' of any fifth point p can be uniquely determined (cf. Theorem 1). \square

We have, thus, a linear algorithm for obtaining H from an outline conic in one view (represented by E')

and four corresponding points across both views. Note that the use of the opacity constraint, via Lemma 2, is less obvious than in the case of reconstruction from point correspondences only. In the case of points (Theorem 1) the opacity constraint was not needed for recovering H , simply because a quadric is uniquely defined by nine points and Lemma 1 provided a simple means for reconstructing the projective coordinates of those nine points. The opacity constraint is needed later only to determine which of the two possible intersection points of Q with the line of sight projects onto the second view. One could trivially extend the case of points to the case of conic and points by first reconstructing a conic of Q from two projections — a quadric is uniquely determined by a conic and four points. However, this is not what is done here.

For practical reasons, it would not be desirable to rely on observing a conic section in both views as this would significantly reduce the generality of our results. In other words, the basic axiom that an object can be represented as a cloud of points would need to be restricted by the additional requirement that some of those points should lie on a conic section in space — not to mention that we would have to somehow identify which of the points lie on a conic section.

As an alternative, we observe the projection of the rim in one of the views and derive the equations for reconstructing the quadric from its outline and four corresponding points. In this case we do not have a conic of Q and four points, and thus it is not a priori clear that a unique reconstruction is possible. Indeed, Theorem 4 shows that without the opacity constraint we have at most eight solutions (16 modulo a common sign flip). This follows from the indeterminacy of whether the conic, projecting onto ψ_1 , is in front or behind (with respect to O) each of the four points. Since the conic in question is the rim, under the opacity assumption the rim is either behind all the points or in front of all the points. Since these two situations differ by a reflection, they correspond to the same quadric (i.e., H up to scale). Thus, the opacity constraint is used here twice — first to recover the quadric’s representation H , and second to determine later (as in Theorem 1) which of the two intersections with the line of sight projects onto the second view ψ_2 . Finally, note that this could have worked only with the rim, and not with any other conic of Q — unless we observe it in both views.

5 Application to Correspondence

In the previous sections we developed the tools for recovering a unique mapping between two projections of an opaque quadric surface. In this section we derive an application of Theorems 1 and 4 to the problem of achieving full correspondence between two grey-level images of a general 3D object.

5.1 Algorithm Using Points Only

For the task of visual correspondence the mapping between two views of a quadric surface will constitute the “nominal quadratic transformation” which, in the case

of points, can be formalized as a corollary of Theorem 1 as follows:

Corollary 2 (of Theorem 1) *A virtual quadric surface can be fitted through any 3D surface, not necessarily a quadric surface, by observing nine corresponding points across two views of the object.*

Proof. It is known that there is a unique quadric surface through any nine points in general position. This follows from a *Veronese* map of degree two, $v_2 : \mathcal{P}^n \rightarrow \mathcal{P}^{(n+1)(n+2)/2-1}$, defined by $(x_0, \dots, x_n) \mapsto (\dots, x^I, \dots)$, where x^I ranges over all monomials of degree two in x_0, \dots, x_n . For $n = 3$, this is a mapping from \mathcal{P}^3 to \mathcal{P}^9 taking hypersurfaces of degree two in \mathcal{P}^3 (i.e., quadric surfaces) into hyperplane sections of \mathcal{P}^9 . Thus, the subset of quadric surfaces passing through a given point in \mathcal{P}^3 is a hyperplane in \mathcal{P}^9 , and since any nine hyperplanes in \mathcal{P}^9 must have a common intersection, there exists a quadric surface through any given nine points. If the points are in general position this quadric is smooth (i.e., H is of full rank).

Therefore, by selecting any nine corresponding points (barring singular configurations) across the two views we can apply the construction described in Theorem 1 and represent the displacement between corresponding points p and p' across the two views as follows:

$$p' \cong (Ap + k_q v') + k_r v', \quad (3)$$

where $k = k_q + k_r$. Moreover, k_q is the relative affine structure of the virtual quadric and k_r is the remaining parallax which we call the residual. The term within parentheses is the nominal quadratic transformation, and the remaining term $k_r v'$ is the unknown displacement along the known direction of the epipolar line. Therefore, Equation 3 is the result of representing the relative affine structure of a 3D object with respect to some reference quadric surface, namely, k_r is a relative affine invariant (because k and k_q are both invariants by Lemma 1). \square

Note that the corollary is analogous to describing shape with respect to a reference plane [9, 26, 28] — instead of a plane we use a quadric and use the tools resulting from Theorem 1 in order to establish a quadric reference surface. The overall algorithm for achieving full correspondence given nine corresponding points p_j, p'_j , $j = 0, 1, \dots, 8$, is summarized below:

1. Determine the epipoles v, v' . This can be done using eight corresponding points to first determine the “fundamental” matrix F satisfying $p'_j{}^\top F p_j = 0$, $j = 1, \dots, 8$. The epipoles follow by $Fv = 0$ and $F^\top v' = 0$ (cf. [4, 5, 27, 28]).
2. Recover the homography A from the equations $Ap_j \cong p'_j$, $j = 1, 2, 3$, and $Av \cong v'$ [27, 21]. This leads to a linear system of eight equations for solving for A up to a scale. Scale v' to satisfy $p'_o \cong Ap_o + v'$.
3. Compute k_j , $j = 4, \dots, 8$ from the equation $p'_j \cong Ap_j + k_j v'$. A least-squares solution is given by the following formula:

$$k_j = \frac{(p'_j \times v')^\top (Ap_j \times p'_j)}{\|p'_j \times v'\|^2}.$$

4. Compute the quadric parameters from the nine equations

$$\begin{pmatrix} x_j \\ y_j \\ 1 \\ k_j \end{pmatrix}^\top H \begin{pmatrix} x_j \\ y_j \\ 1 \\ k_j \end{pmatrix} = 0, \quad (4)$$

for $j = 0, 1, \dots, 8$. Note that $k_o = 1$ and $k_1 = k_2 = k_3 = 0$.

5. For every other point p compute k_q as the appropriate root of k of $ak^2 + b(p)k + c(p) = 0$, where the coefficients a, b, c follow from $(x_q, y_q, 1, k_q)H(x_q, y_q, 1, k_q)^\top = 0$, and the appropriate root follows from the sign of r for $ak_o^2 + b(p_o)k_o + c(p_o) = 0$ consistent with the root $k_o = 1$.
6. Warp ψ_1 according to the nominal transformation

$$\bar{p} \cong Ap + k_q \mathbf{v}'.$$

Thus, the image brightness at any $p \in \psi_1$ is copied onto the transformed location \bar{p} .

7. The remaining displacement (residual) between p' and \bar{p} consists of an unknown displacement k_r along the known epipolar line:

$$p' \cong \bar{p} + k_r \mathbf{v}'.$$

The spatio-temporal derivatives of image brightness can be used to recover k_r .

This algorithm was implemented and applied to the pair of images displayed in the top row of Fig. 2. Note that typical displacements between corresponding points around the center region of the face vary around 20 pixels. Achieving full correspondence between two views of a face is challenging for two reasons. First, a face is a complex object which is not easily parameterized. Second, the texture of a typical face does not contain enough image structure to obtain point-to-point correspondence in a reliable manner. However, there are a few points (on the order of 10–20) that can be reliably matched, such as the corners of the eye, mouth and eyebrows. We rely on these few points to recover the epipolar geometry and the nominal quadratic transformation.

Fig. 3 displays the results in the following manner. The top row display shows the original second view. Notice that the transformed first view (middle row display) appears to be heading in the right direction but is slightly deformed. The selection of corresponding points (selected manually) yielded an ellipsoid whose outline on the first view circumscribes the image of the head (this is not a general phenomenon; see later in this section). The overlay between the edges of the original second view and the edges of the transformed view are also shown in the middle row display. Notice that the residuals are relatively small, typically in the range of 1–2 pixels. The residuals are subsequently recovered by using a coarse-to-fine gradient-based optical flow method following [11, 3] constrained along epipolar directions (cf. [24]). The final results are shown in the bottom row display.

Also, a tight fit of a quadric surface onto the object can be obtained by using many corresponding points to

obtain a least-squares solution for H . Note that from a practical point of view we would like the quadric to lie as close as possible to the object; otherwise the algorithm, though correct, would not be useful, i.e., the residuals may be larger than the original displacements between the two views. In this regard, the re-parameterization suggested in Theorem 2 may provide a better fit for least-squares methods. Using the parameterization described in the theorem, the entries h_{11}, h_{22}, h_{33} vanish, leaving only six parameters of H to be determined (see proof of Theorem 2). Thus, instead of recovering nine parameters in a least-squares solution, we solve for only six parameters, which is equivalent to constraining the resulting quadric to lie on three object points. The implementation steps described above should be modified in a way that readily follows from the proof of Theorem 2.

We have seen that the quadric’s outline in the example shown in Fig. 3 circumscribes the image of the object. This, however, is not a general property and the issue is taken further in the next section where the results of Theorem 4 become relevant and practical.

5.2 Algorithm Using Conic and Points

When only point matches are used, one cannot guarantee that the outline of the recovered quadric will circumscribe the image of the object. Some choice of corresponding points may give rise to a quadric whose outline happens to fall within the image of the object. Fig. 4 illustrates this possibility on a different face-pair. One can see that the outline of the quadric (again an ellipsoid) encompasses all sample points, but inscribes the image of the head, leaving out the peripheral region.

In general, points p outside the outline correspond to rays \overline{OP} that do not intersect the quadric in real space, and therefore the corresponding k_q are complex conjugate (i.e., the nominal quadratic transformation cannot be applied to p). This is where Theorem 4 becomes useful in practice. We have shown there that instead of nine corresponding points, the outline of the quadric and four corresponding points are sufficient for uniquely determining the mapping between the two views due to the quadric. In the context of visual correspondence, the outline conic can be set arbitrarily (such as circumscribing the image of the object of interest), and the rest follows from Theorem 4. This is formalized in the following corollary:

Corollary 3 (of Theorem 4) *A virtual quadric surface lying on four object points projecting onto an arbitrary outline (conic) can be fitted through any 3D surface, not necessarily a quadric surface, by observing the corresponding four point matches across two views of the object.*

The algorithm for recovering a virtual quadric reference surface, by setting an arbitrary conic in the first view, is summarized below. We are given four corresponding points p_j, p'_j , $j = 0, 1, 2, 3$ and the epipoles v, v' . The homography A due to the plane of reference passing through P_j , $j = 1, 2, 3$, is recovered as before (steps 1 and 2 of the point-based algorithm). The rest goes as follows:

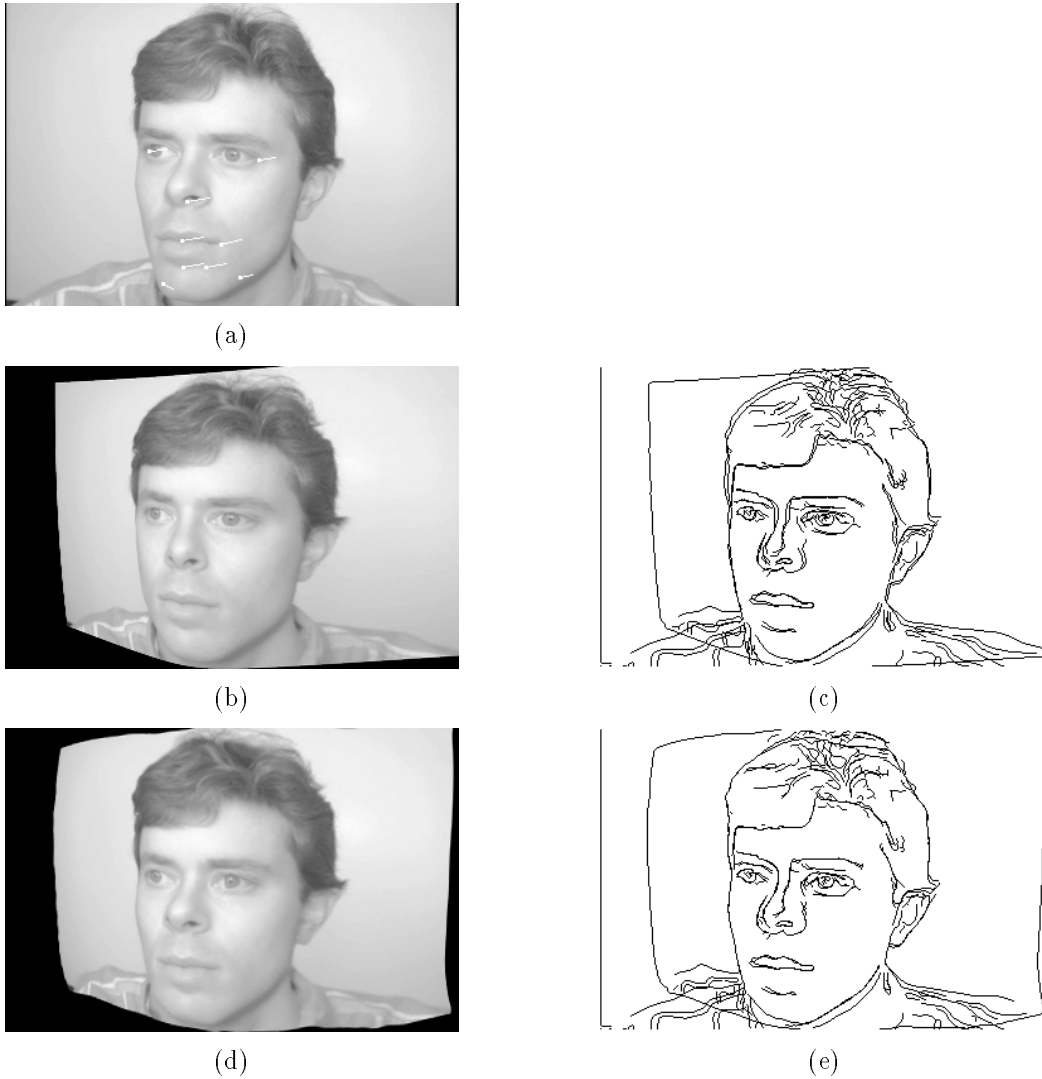


Figure 3: Nominal quadratic transformation from nine corresponding points and subsequent refinement of residual displacements using optical flow. (a) Original second view ψ_2 (the first view, ψ_1 , is shown in Fig. 2). Nine corresponding points were manually chosen. The needle heads mark the positions of the sampling points in ψ_2 and the needles denote the corresponding displacement vectors. (b) The view ψ_1 warped using the nominal quadratic transformation. (c) The residual displacement shown by overlaying the edges of (a) and (b). Note that the typical displacements are within 1–2 pixels (the original displacements were in the range of 20 pixels; see Fig. 2). (d) The image in (b) warped further by applying optic flow along epipolar lines towards ψ_2 . (e) The performance of the correspondence strategy (nominal transformation followed by optical flow) is illustrated by overlaying the edges of (a) and (d). Note that correspondence has been achieved to within subpixel accuracy almost everywhere.

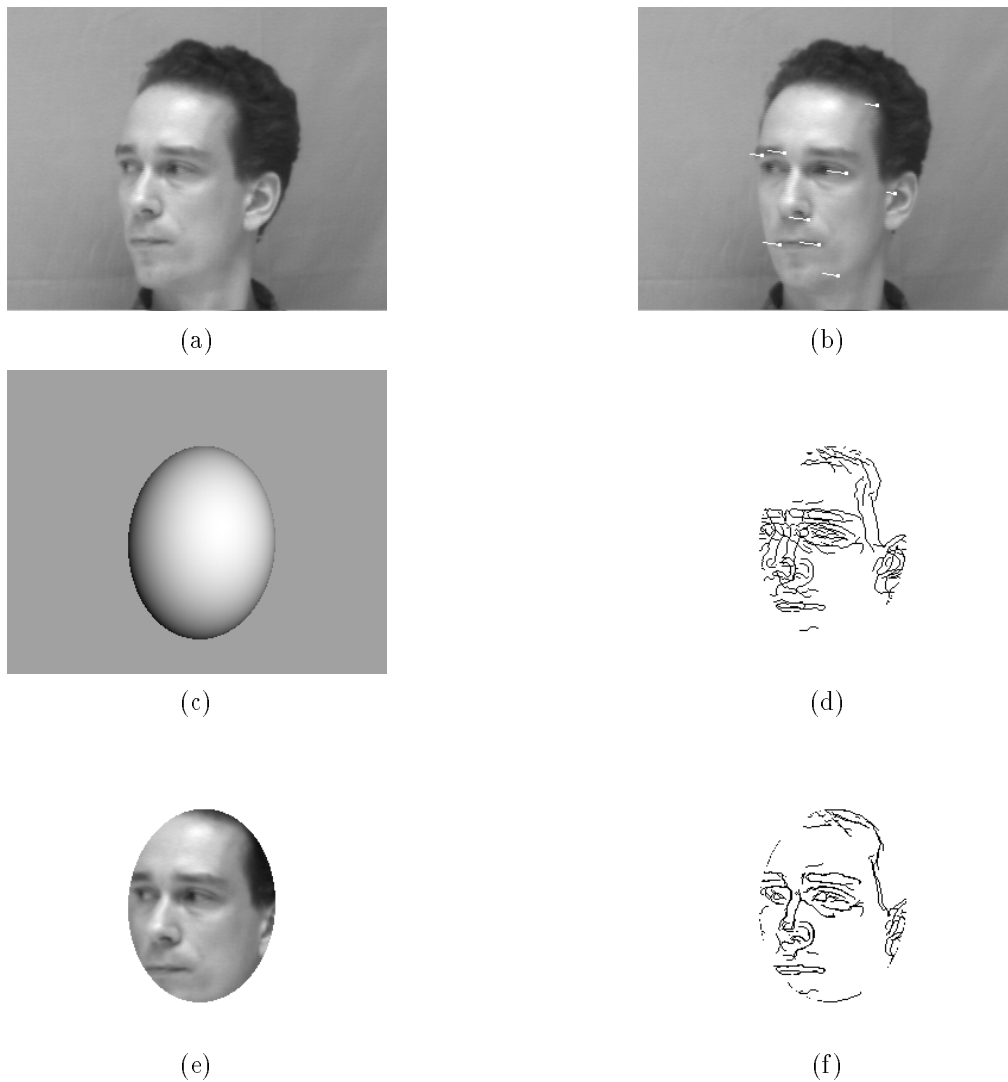


Figure 4: Nominal quadratic transformation from nine corresponding points: a case where the quadric's outline inscribes the image of the object. (a) ψ_1 , (b) ψ_2 with overlaid corresponding points used for recovering the nominal quadratic transformation. (c) The recovered quadric (values of k_q). The uniform grey background indicates complex conjugate values for the roots. (d) The overlaid edges of ψ_1 and ψ_2 masked by the ellipse having real roots. (e) The masked region of the transformed first view. (f) The overlaid edges of ψ_2 (b) and the transformed view (e). Note that within the masked area of real k_q -values, the residuals are fairly small.



Figure 5: Nominal quadratic transformation recovered from a conic and four corresponding points. (a) original view ψ_2 with corresponding points overlaid. (b) The transformed first view, ψ_1 , within the given conic (circle around the face in this example).



Figure 6: The nominal quadric transformation due to a hyperboloid of two sheets. This unintuitive solution due to a deliberately unsuccessful choice of sample points creates the mirror image on the right side that is due to the second sheet of the hyperboloid.

1. Select an arbitrary conic $p^\top E' p = 0$ (presumably one that circumscribes the image of the object in the first view).
2. Solve for vector h and scalar h_{44} from the system

$$p_j^\top h + h_{44} k_j = \sqrt{p_j^\top E' p_j},$$

$j = 0, 1, 2, 3$. Note that $k_0 = 1$ and $k_1 = k_2 = k_3 = 0$.

3. The parameter matrix H representing the quadric Q is given by

$$H = \begin{pmatrix} hh^\top - E' & h_{44}h \\ h_{44}h^\top & h_{44}^2 \end{pmatrix}.$$

The remaining steps are the same as steps 5, 6, and 7 in the point-based algorithm.

This algorithm was also implemented and applied to the pair of images used earlier (top row of Fig. 2). The arbitrary conic was chosen to be a circle circumscribing the image of the head in one view. Fig. 5 shows the original second view and the warped first view according to the recovered quadratic nominal transformation due to the conic and only four corresponding points.

Finally, although ellipsoids and paraboloids are the most natural quadric surfaces for this application, we cannot (in principle) eliminate other classes of quadrics from appearing in this framework. For example, a hyperboloid of two sheets may yield unintuitive results, under specialized circumstances (see Fig. 6). Since the recovered quadrics are in real space, a certain limited classification is possible (based on the ranks of the matrices and the sign pattern of the eigenvalues of H), but unfortunately that classification is not sufficient to eliminate hyperboloids of two sheets. In practice, however, the situation illustrated in Fig. 6 is accidental and was contrived for purposes of illustrating this kind of failure mode.

6 Discussion

The theoretical part of this paper addressed the question of establishing a one-to-one mapping between two views of an unknown quadric surface. We have shown

that nine corresponding points are sufficient to obtain a unique map, provided we make the assumption that the surface is opaque. Similarly, four corresponding points and the outline conic of the quadric in one view are sufficient to obtain a unique map as well. We have also shown that an appropriate parameterization of the image planes facilitates certain questions of interest such as the likelihood that eight corresponding points will be coming from a quadric lying in the vicinity of both centers of projection.

On the practical side, we have shown that the tools developed for quadrics can be applied to any 3D object by setting up a virtual quadric surface lying in the vicinity of the object. The quadric serves as a reference surface, but also facilitates the correspondence problem. For example, given the epipoles (which can be recovered independently), by specifying a conic circumscribing the image of an object in one view and observing four corresponding points with the other view one can obtain the virtual quadric surface whose rim projects to the specified outline conic and which lies on the four corresponding object points in 3D space. The virtual quadric induces a unique mapping between the two views (the nominal quadratic transformation), which is equivalent to projecting the object onto the quadric along the projection lines toward the first view, followed by a projection of the quadric onto the second view. What remains are residual displacements along epipolar lines whose magnitude are small in regions where the object lies close to the virtual quadric. The residual displacements are later refined by use of local spatio-temporal detectors that implement the constant brightness equation, or any correlation scheme (cf. [14, 22, 1]), along the epipolar lines. In the implementation section we have shown that two views of a face with typical displacements of around 20 pixels are brought closer to displacements of around 1–2 pixels by the transformation. Most optical flow methods can deal with such small displacements quite effectively.

On the conceptual level, two proposals were made. First, the correspondence problem is treated as a two-stage process combining geometric information captured by a small number of point matches, and photometric information captured by the spatio-temporal derivatives of image brightness. Second, manipulations on 3D object space are achieved by first manipulating a reference surface. The reference surface is viewed here as an approximate prototype of the observed object, and shape is measured relative to the prototype rather than relative to a generic (minimal) frame of reference.

The notion of reference surfaces as prototypes may be relevant for visual recognition, visual motion and stereopsis. In some of these areas one may find some support to this notion in the human vision literature, although not directly. For example, the phenomenon of “motion capture” introduced by Ramachandran [18, 19, 20] is suggestive of the kind of motion measurement presented here. Ramachandran and his collaborators observed that the motion of certain salient image features (such as gratings or illusory squares) tends to dominate the perceived motion in the enclosed area by masking incoherent mo-

tion signals derived from uncorrelated random dot patterns, in a winner-take-all fashion. Ramachandran therefore suggested that motion is computed by using salient features that are matched unambiguously and that the visual system assumes that the incoherent signals have moved together with those salient features [18]. The scheme suggested in this paper may be considered as a refinement of this idea. Motion is “captured” in Ramachandran’s sense by the reference surface, not by assuming the motion of the salient features but by computing the nominal motion transformation. The nominal motion is only a first approximation which is further refined by use of spatio-temporal detectors, provided that the remaining residual displacement is in their range, namely, the object being tracked and the reference surface model are sufficiently close. In this view the effect of capture attenuates with increasing depth of points from the reference surface, and is not affected, in principle, by the proximity of points to the salient features in the image plane.

Other suggestive data include stereoscopic interpolation experiments by Mitchison and McKee [16]. They describe a stereogram which has a central periodic region bounded by unambiguously matched edges. Under certain conditions the edges impose one of the expected discrete matchings (similar to stereoscopic capture; see also [17]). Under other conditions a linear interpolation in depth occurs between the edges violating any possible point-to-point match between the periodic regions. The linear interpolation in depth corresponds to a plane passing through the unambiguously matched points, which supports the idea that correspondence starts with the computation of nominal motion (in this case due to a planar reference surface), determined by a small number of salient unambiguously matched points, and is later refined using short-range mechanisms.

To conclude, the computational results provide tools for further exploring the utility of reference surfaces in visual applications, and provide specific applications to the task of visual correspondence (visual motion).

Acknowledgments

Thanks to Azriel Rosenfeld for critical reading and comments on the final draft of this manuscript; to Tomaso Poggio for helpful discussions on visual correspondence. Also thanks to David Beymer for providing some of the images used for our experiments, and to Long Quan for providing the code we used for recovering epipoles. A. Shashua is supported by a McDonnell-Pew postdoctoral fellowship from the Department of Brain and Cognitive Sciences. S. Toelg was supported by a postdoctoral fellowship from the Deutsche Forschungsgemeinschaft while he was at MIT. Part of this work was done while S. Toelg was at the Institut fuer Neuroinformatik, Ruhr-Universitaet Bochum, Germany.

References

- [1] E.H. Adelson and J.R. Bergen. Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America*, 2:284–299, 1985.
- [2] J.R. Bergen, P. Anandan, K.J. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. In *Proceedings of the European Conference on Computer Vision*, Santa Margherita Ligure, Italy, June 1992.
- [3] J.R. Bergen and R. Hingorani. Hierarchical motion-based frame rate conversion. Technical report, David Sarnoff Research Center, 1990.
- [4] O.D. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In *Proceedings of the European Conference on Computer Vision*, pages 563–578, Santa Margherita Ligure, Italy, June 1992.
- [5] O.D. Faugeras, Q.T. Luong, and S.J. Maybank. Camera self calibration: Theory and experiments. In *Proceedings of the European Conference on Computer Vision*, pages 321–334, Santa Margherita Ligure, Italy, June 1992.
- [6] B.K.P. Horn. Relative orientation. *International Journal of Computer Vision*, 4:59–78, 1990.
- [7] M. Irani, B. Rousso, and S. Peleg. Robust recovery of ego-motion. In D. Chetverikov and W. Kropatsch, editors, *Computer Analysis of Images and Patterns (Proc. of CAIP’93)*, pages 371–378, Budapest, Hungary, September 1993. Springer.
- [8] D. W. Jacobs. Generalizing invariants for 3-D to 2-D matching. In *Proceedings of the 2nd European Workshop on Invariants*, Ponta Delagada, Azores, October 1993.
- [9] J.J. Koenderink and A.J. Van Doorn. Affine structure from motion. *Journal of the Optical Society of America*, 8:377–385, 1991.
- [10] R. Kumar and P. Anandan. Direct recovery of shape from multiple views: A parallax based approach. In *Proceedings of the International Conference on Pattern Recognition*, Jerusalem, Israel, October 1994.
- [11] B.D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings IJCAI*, pages 674–679, Vancouver, Canada, 1981.
- [12] Q.T. Luong, R. Deriche, O.D. Faugeras, and T. Papadopoulos. On determining the fundamental matrix: Analysis of different methods and experimental results. Technical Report INRIA, France, 1993.
- [13] H. A. Mallot, H. H. Bülthoff, J. J. Little, and S. Bohrer. Inverse perspective mapping simplifies optical flow computation and obstacle detection. *Biological Cybernetics*, 64:177–185, 1991.
- [14] D. Marr and S. Ullman. Directional selectivity and its use in early visual processing. *Proceedings of the Royal Society of London B*, 211:151–180, 1981.
- [15] S.J. Maybank. The projective geometry of ambiguous surfaces. *Proceedings of the Royal Society of London*, 332:1–47, 1990.
- [16] G.J. Mitchison and S.P. McKee. Interpolation in stereoscopic matching. *Nature*, 315:402–404, 1985.

- [17] K. Prazdny. ‘Capture’ of stereopsis by illusory contours. *Nature*, 324:393, 1986.
- [18] V.S. Ramachandran. Capture of stereopsis and apparent motion by illusory contours. *Perception and Psychophysics*, 39:361–373, 1986.
- [19] V.S. Ramachandran and P. Cavanagh. Subjective contours capture stereopsis. *Nature*, 317:527–530, 1985.
- [20] V.S. Ramachandran and V. Inada. Spatial phase and frequency in motion capture of random-dot patterns. *Spatial Vision*, 1:57–67, 1985.
- [21] L. Robert and O.D. Faugeras. Relative 3D positioning and 3D convex hull computation from a weakly calibrated stereo pair. In *Proceedings of the International Conference on Computer Vision*, pages 540–544, Berlin, Germany, May 1993.
- [22] J.P.H. Van Santen and G. Sperling. Elaborated Reichardt detectors. *Journal of the Optical Society of America*, 2:300–321, 1985.
- [23] H.S. Sawhney. 3D geometry from planar parallax. Technical report, IBM Almaden Research Center, April 1994.
- [24] A. Shashua. Correspondence and affine shape from two orthographic views: Motion and Recognition. A.I. Memo No. 1327, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, December 1991.
- [25] A. Shashua. *Geometry and Photometry in 3D Visual Recognition*. PhD thesis, M.I.T Artificial Intelligence Laboratory, AI-TR-1401, November 1992.
- [26] A. Shashua. On geometric and algebraic aspects of 3D affine and projective structures from perspective 2D views. In *Proceedings of the 2nd European Workshop on Invariants*, Ponta Delagada, Azores, October 1993. Also MIT AI Memo No. 1405, July 1993.
- [27] A. Shashua. Projective structure from uncalibrated images: structure from motion and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1994. In press.
- [28] A. Shashua and N. Navab. Relative affine structure: Theory and application to 3D reconstruction from perspective views. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Seattle, Washington, 1994.
- [29] K. Storjohann, Th. Zielke, H. A. Mallot, and W. von Seelen. Visual obstacle detection for automatically guided vehicles. In *Proceedings of IEEE Conference on Robotics and Automation*, pages 761–766, Los Alamitos, CA, 1990.
- [30] Y. Zheng, D. G. Jones, S. A. Billings, J. E. W. Mayhew, and J. P. Frisby. Switcher: A stereo algorithm for ground plane obstacle detection. *Image and Vision Computing*, 8:57–62, 1990.