Quadric Reconstruction from Dual-Space Geometry

Geoffrey Cross and Andrew Zisserman
Robotics Research Group, Department of Engineering Science
University of Oxford, Oxford, OX1 3PJ.

Abstract

We describe the recovery of a quadric surface from its image in two or more perspective views. The recovered quadric is used in 3D modeling and image registration applications.

There are three novel contributions. First, it is shown that a one parameter family of quadrics is recovered from outlines in two views. The ambiguity is reduced to two-fold given a point correspondence. There is no ambiguity from outlines in three or more views. Second, it is shown that degenerate quadrics reduce the ambiguity of reconstruction. Third, it is shown that surfaces can be piecewise quadric approximated from piecewise conic approximations of their outlines.

All these cases are illustrated by examples with real images. Implementation details are given and the quality of the results is assessed.

1 Introduction

The aim of the work described here is to recover quadric surfaces (spheres, cylinders, ...) in 3D from two or more perspective images. The quadrics can be recovered from points on their surface, or from their (conic) outlines, or from a combination of the two. We are using this quadric recovery in two applications: 3D surface recovery, where 3D VRML models are built automatically from uncalibrated image sequences [3]; and image registration, where images of curved surfaces are registered in order to obtain super-resolution surface textures from image sequences [1, 8, 13].

The automatic recovery of explicit quadric surfaces from multiple views is a significant improvement over the common alternative of triangulating the smooth surface. Here we build on the work of Shashua and Toelg [16], Ma *et al.* [11, 12], Quan [14], and Karl *et al.* [9]. The novelty of the work in this paper is that the reconstruction ambiguity is defined precisely using both linear algebra and geometric arguments. The vehicle for this is the tangent dual space. All the cases have been implemented on real images using automatic robust processing and optimal fitting.

It is assumed that the multiple view relations are available *a priori*. These include the fundamental matrix for two views, and the trifocal tensor for three. These tensors are computed automatically from point (corner) correspondences [2, 17, 19]. Consequently, a consistent set of camera projection matrices are available, so that 3D structure may be recovered up to a homography of 3-space.

2 Quadrics — Review

A (point) quadric is a surface represented by a 4×4 symmetric matrix $\mathbf{Q}$, such that points $\mathbf{X}$ on the surface satisfy the quadratic constraint,

$$\mathbf{X}^\top \mathbf{Q} \mathbf{X} = 0, \tag{1}$$

where $\mathbf{X}$ is a homogeneous 4-vector which represents a point in 3D, $\mathbf{X} = (x, y, z, 1)^\top$. A quadric has nine degrees of freedom corresponding to the independent elements of $\mathbf{Q}$ up to an overall scale.

Recovery from point correspondences Equation (1) can be rewritten in the form

$$\mathtt{a}\mathbf{v} = 0 \ ,$$

where $\mathtt{a}$ is a 1 by 10 matrix which is a function of the point $\mathbf{X}$ only, whilst $\mathbf{v}$ is an homogeneous 10-vector containing the distinct matrix elements of $\mathbf{Q}$. Each 3D point provides a similar constraint, so that from 9 points a matrix equation $\mathtt{A}\mathbf{v} = \mathbf{0}$ can be constructed, where $\mathtt{A}$ is 9×10 matrix formed from the stacked matrices $\mathtt{a}$. Provided the nine points are in general position the quadric is obtained as the one dimensional null-space of $\mathtt{A}$. This shows that a quadric is uniquely defined by 9 points in general position [15].

Suppose 3D points $\mathbf{X}$ are imaged as $\mathbf{x} = \mathtt{P}\mathbf{X}$, $\mathbf{x}' = \mathtt{P}\mathbf{X}$ in an image pair, where the image points $\mathbf{x}, \mathbf{x}'$ are represented by homogeneous 3-vectors, $\mathbf{x} = (x, y, 1)^\top$, and $\mathtt{P}, \mathtt{P}'$ are 3×4 camera projection matrices for the first and second view. Given the fundamental matrix $\mathtt{F}$ for the view pair, then from [10] the camera matrices can be chosen as,

$$\mathtt{P} = [\mathtt{I} \mid \mathbf{0}], \quad \mathtt{P}' = [[\mathbf{e}']_\times \mathtt{F} \mid \mathbf{e}'] \ ,$$

where $\mathbf{e}'$ is the epipole in the second image ($\mathtt{F}^\top \mathbf{e}' = \mathbf{0}$) and $[\mathbf{e}']_\times$ is the 3×3 skew matrix such that $[\mathbf{e}']_\times \mathbf{x} =$

$\mathbf{e}' \times \mathbf{x}$. A 3D point $\mathbf{X}$ is reconstructed from its image correspondence $\mathbf{x} \leftrightarrow \mathbf{x}'$ by back-projection (via $\mathtt{P}, \mathtt{P}'$) and triangulation [7].

It follows that a quadric can be uniquely reconstructed (in the projective frame defined by $\mathtt{P}, \mathtt{P}'$) from nine image point correspondences [16]. If there are only eight 3D points (from 8 correspondences) then $\mathtt{A}$ has a two dimensional null-space, and there is a one parameter family (a pencil) of quadrics passing through the points.

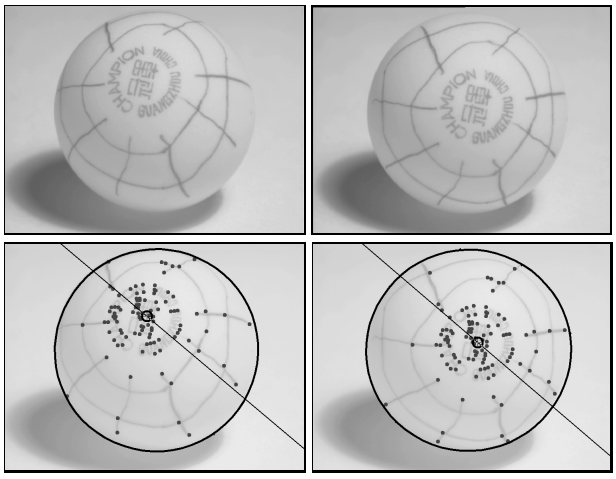


Figure 1: Two images of a ping-pong ball — an example of a quadric surface. The original images are shown above, and the lower image pair shows the corner matches used in reconstruction. Also shown are the conic outlines computed from the fitted quadric. Note the good agreement with the actual ball outlines. The circled point in the left image is correctly mapped to the circled point in the right image by the quadric induced mapping. The epipolar lines for the circled points are also shown.

Quadric induced transfer It is well known that a world plane induces a mapping between images (a plane projectivity) such that the images of 3D points on the plane are mapped to each other. Similarly, a quadric induces a mapping between images such that images of 3D points on the quadric surface are mapped to each other.

A point $\mathbf{x}$ in one view back-projects to a ray which intersects the quadric in two points (these points may be real, coincident or complex). These two intersection points are then imaged in the second view, to produce two candidate $\mathbf{x}'$s. Coincident points occur when the ray is tangent to the quadric, in which case $\mathbf{x}$ is on the outline in the first view. Real and complex non-coincident intersections occur according to whether the pre-image of $\mathbf{x}$ is on the quadric or not.

It can be shown [6] that the mapping may be written $\mathbf{x}' = \lambda_{\pm}[\mathbf{e}']_{\times}\mathbf{F}\mathbf{x} + \mathbf{e}'$, where

$$\lambda_{\pm} = \left(-\mathbf{q}^{\top}\mathbf{x} \pm \sqrt{(\mathbf{q}^{\top}\mathbf{x})^2 - (\mathbf{x}^{\top}\mathbf{Q}_{33}\mathbf{x})q}\right) / (\mathbf{x}^{\top}\mathbf{Q}_{33}\mathbf{x})$$

and $\mathbf{Q}$ is decomposed as

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{33} & \mathbf{q} \\ \mathbf{q}^{\top} & q \end{bmatrix} \tag{2}$$

with $\mathbf{Q}_{33}$ a symmetric 3 by 3 matrix, $\mathbf{q}$ a 3 vector, and q a scalar. A transfer example is shown in figure 1.

Quadric outlines The outline of a quadric in a perspective image is a conic. A (point) conic is represented by a 3×3 symmetric matrix $\mathbf{C}$, such that points $\mathbf{x}$ on the conic satisfy the quadratic constraint,

$$\mathbf{x}^{\top}\mathbf{C}\mathbf{x} = 0 \ . \tag{3}$$

A conic has five degrees of freedom corresponding to the independent elements of $\mathbf{C}$ up to an overall scale.

The conic outline of $\mathbf{Q}$ in the first image (with $\mathtt{P} = [\mathtt{I} \mid \mathbf{0}]$) is

$$\mathbf{C} = \left(\mathbf{q}\mathbf{q}^{\top} - q\mathbf{Q}_{33}\right) \ . \tag{4}$$

A similar, but more complicated, expression is obtained for $\mathtt{P}'$.

A simple expression is obtained for any $\mathtt{P}$ by using *duals*. The quadric dual is defined by $\mathbf{Q}^* = \mathrm{adj}(\mathbf{Q})$, where $\mathrm{adj}(\mathbf{Q})$ is the adjoint of $\mathbf{Q}$. $\mathbf{Q}^*$ is the tangent dual of $\mathbf{Q}$ such that if $\boldsymbol{\pi}$ is a plane tangent to $\mathbf{Q}$, then $\boldsymbol{\pi}^{\top}\mathbf{Q}\boldsymbol{\pi} = 0$, with $\boldsymbol{\pi}$ an homogeneous 4-vector. It can be shown that [9, 12, 18],

$$\mathbf{C}^* = \mathtt{P}\mathbf{Q}^*\mathtt{P}^{\top} \ , \tag{5}$$

where $\mathbf{C}^* = \mathrm{adj}(\mathbf{C})$ is the conic dual [15]. This relation between the dual quadric and dual conic outline shows that the outline places at most 5 constraints on the quadric, corresponding to the 5 degrees of freedom of the conic. An example of a computed outline is shown in figure 1.

3 Reconstruction from outlines in two views

3.1 Geometric discussion

Figure 2a shows a cone tangent to the quadric with the camera centre as vertex. This cone intersects the image plane in the conic outline. The single view relation on duals (5) shows that, loosely, the cone is the equivalent of 5 tangent planes corresponding to

the 5 degrees of freedom of the conic. A second cone from the second view (figure 2b) also imposes 5 tangent plane constraints. However, two planes are in common (see figure 2c), so that the net number of tangent plane constraints is eight. Since tangency to nine planes uniquely define a quadric [15], there is one degree of freedom remaining for the quadric, i.e. two outlines determine a pencil of dual quadrics.

In the image plane the common quadric tangent planes intersect the image in epipolar lines tangent to the conic outlines, i.e. the epipolar tangents to the conics correspond. This is shown in figure 3. A consequence of this *consistency* constraint is that although the two conics have 10 degrees of freedom in total, two are lost as they must conform to the epipolar geometry. Thus, two image conics only give 8 constraints on the quadric.

This can be further illustrated in the dual-space. Figure 2d shows the dual quadric, and it's intersection with the planes which are dual to the camera centres. Each conic of intersection is dual to the real-space viewing cone, and the line of intersection of the two planes is dual to the base-line in real-space. The dual quadric must pass through each of these dual conics. It can be shown [6] that there is a pencil of quadrics satisfying these constraints.

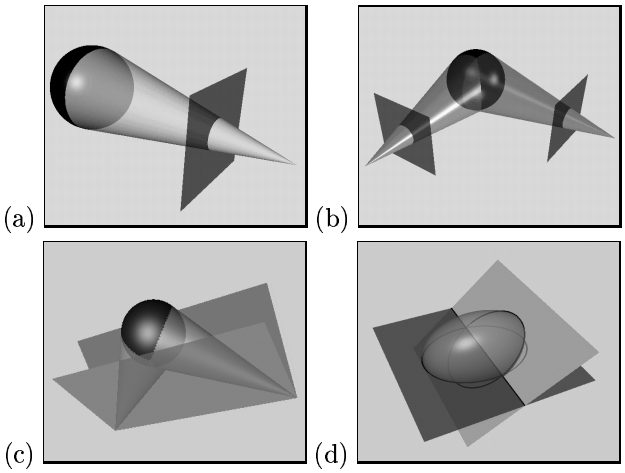


Figure 2: (a) The conic outline of a quadric back-projects to a viewing cone tangent to the quadric with vertex the camera centre. (b) Each view generates a cone. (c) The epipolar tangency constraint forces the two cones to share a pair of planes. (d) The dual-space equivalent of two cameras viewing a single quadric. The planes are dual to the camera centres, whilst the conics are dual to the viewing cones.

3.2 Algebraic derivation

The reconstruction ambiguity is now derived using linear algebra. Equation (5) is a linear relationship between the dual quadric and the dual image conic. For two cameras, with projection matrices, $\mathtt{P}$ and $\mathtt{P}'$,

$$\begin{aligned} \mathbf{C}^* &= \mathtt{P}\mathbf{Q}^*\mathtt{P}^\top \\ \mathbf{C}'^* &= \mathtt{P}'\mathbf{Q}^*\mathtt{P}'^\top \ . \end{aligned}$$

Note, that since the quadric is imaged using projection matrices consistent with the epipolar geometry (since $\mathtt{P}, \mathtt{P}'$ are produced from $\mathtt{F}$ as described in section 2) the conics are guaranteed to satisfy the epipolar tangent consistency constraint.

Each of these equations can be rewritten in the form

$$\begin{bmatrix} \mathbf{C}^*_{11} \\ \vdots \\ \mathbf{C}^*_{33} \end{bmatrix} = \mathtt{B} \begin{bmatrix} \mathbf{Q}^*_{11} \\ \vdots \\ \mathbf{Q}^*_{44} \end{bmatrix} \tag{6}$$

or more concisely as $\mathbf{c}^* = \mathtt{B}\mathbf{v}^*$ where $\mathtt{B}$ is a 6 by 10 matrix which is quadratic in the elements of the projection matrix $\mathtt{P}$, and $\mathbf{c}^*, \mathbf{v}^*$ are the dual conic and dual quadric written as 6- and 10- vectors respectively. Although there are 6 equations, only 5 of these are independent, as the equality is only up to scale. If the scaling is included explicitly then there are 6 independent equations for 6 unknowns. Then $\mathbf{c}^* = \alpha\,\mathtt{B}\mathbf{v}^*$ and $\mathbf{c}'^* = \beta\,\mathtt{B}'\mathbf{v}^*$, where α and β are the unknown scales. These equations can be written as a single matrix equation

$$\mathtt{M}\mathbf{w} = \mathbf{0} \tag{7}$$

where

$$\mathtt{M} = \begin{bmatrix} \mathtt{B} & \mathbf{c}^* & \mathbf{0} \\ \mathtt{B}' & \mathbf{0} & \mathbf{c}'^* \end{bmatrix} \qquad \mathbf{w} = \begin{bmatrix} \mathbf{v}^* \\ \alpha \\ \beta \end{bmatrix}$$

and $\mathtt{M}$ is a 12 by 12 matrix. However, only 10 of the *joint* equations are independent—this is a consequence of the epipolar geometry consistency constraint on the conics described above. It follows that $\mathtt{M}$ has rank 10 and thus has a two dimensional null space. From this null space $\mathbf{Q}^*$ is determined as $\mathbf{Q}^*(\lambda) = \mathbf{Q}^*_1 + \lambda\mathbf{Q}^*_2$ where $\mathbf{Q}^*_1, \mathbf{Q}^*_2$ are 4×4 matrices corresponding to the 10-vectors which span the $\mathbf{v}^*$ part of the null space of $\mathtt{M}$. $\mathbf{Q}^*(\lambda)$ is a one parameter family, with λ the scalar parameter, as $\mathbf{Q}^*$ is only defined up to scale.

3.3 Reducing the ambiguity

It has been shown that a pair of conic outlines from two views allows the dual quadric to be reconstructed up to a one parameter family in the form

$$\mathbf{Q}^*(\lambda) = \mathbf{Q}^*_1 + \lambda\mathbf{Q}^*_2 \ .$$

If another tangent plane $\boldsymbol{\pi}$ were to be available then $\boldsymbol{\pi}^\top \mathbf{Q}^*(\lambda)\boldsymbol{\pi} = 0$ would determine λ uniquely, and thus $\mathbf{Q}^*$ and the point quadric $\mathbf{Q} = (\mathbf{Q}^*)^*$ uniquely.

A 3D point $\mathbf{X}$ on the quadric reduces the one parameter family to two solutions. To see this note that the adjoint of $\mathbf{Q}^*(\lambda)$ is cubic in λ, and the constraint that $\mathbf{X}$ is on $\mathbf{Q}$, i.e. $\mathbf{X}^\top\mathbf{Q}\mathbf{X} = 0$, produces a cubic in λ. However, one of the three solutions always corresponds to a degenerate quadric (repeated planes). Details are given in [6]. It follows that:

There is a two fold reconstruction ambiguity for a quadric given its outline and the image of one point on its surface in two views.

Any further (general) point on the surface will resolve this ambiguity.

3.4 Implementation details

Point correspondences are first corrected to be consistent with the epipolar geometry using the polynomial algorithm of Hartley and Sturm [7]. This algorithm determines corresponding points $\hat{\mathbf{x}}, \hat{\mathbf{x}}'$ which exactly satisfy the epipolar geometry and minimise the image distances from the measured points $\mathbf{x}, \mathbf{x}'$. The quadric is fitted to the 3D point back projected from $\hat{\mathbf{x}}, \hat{\mathbf{x}}'$.

An analogous correction is required for the measured conic outlines before they are used for estimation in equation (7). As described in section 3.1, to be consistent with epipolar geometry the conic epipolar tangents must correspond. We seek conics $\hat{\mathtt{C}}, \hat{\mathtt{C}}'$ which satisfy the epipolar tangency constraint and are "close to" the measured conics $\mathtt{C}, \mathtt{C}'$. This correction is applied in two stages.

1. Conic outline correction The conic outlines are detected using a Canny edge-detector [5], and conics are fitted to the edgel list using Bookstein's fitting algorithm [4].

The following (non-optimal) algorithm, adjusts $\mathtt{C}'$, the conic in the second image, to be consistent with $\mathtt{C}$, the conic in the first. This is not optimal because $\mathtt{C}$ is assumed without measurement error. This initial correction is used as a starting point for a non-linear estimation described below.

The conic in the second image is defined uniquely by tangency to two lines, the point of tangency on these lines, and another point. The two lines ($\mathbf{l}'_1$ and $\mathbf{l}'_2$) are provided by mapping the epipolar tangents ($\mathbf{l}_1$ and $\mathbf{l}_2$) of $\mathtt{C}$ to the second image using the epipolar line homography. The tangency points are provided by the points on $\mathbf{l}'_1$ and $\mathbf{l}'_2$ closest to the *measured* conic in the second image. The single point is obtained from the *measured* conic. Details are in [6]. In practise the lines ($\mathbf{l}'_1$ and $\mathbf{l}'_2$) are within a pixel of the measured conic, so the correction is minor. See figure 3.

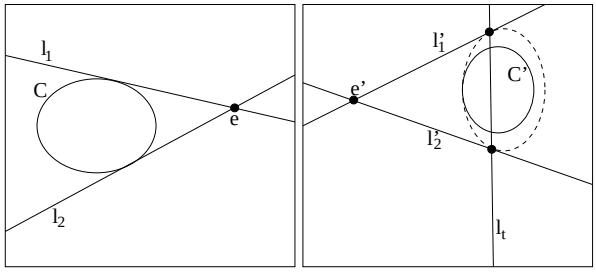


Figure 3: The epipolar geometry imposes an epipolar tangency constraint on the conic outlines of a quadric in two images. The edge-fitted conics are therefore corrected to ensure that they conform to this constraint. The dotted conic is the corrected one (exaggerated error).

2. Optimising the correction The cost function to be minimised is the sum of squared perpendicular distance of the measured edgels from the conic outlines $\hat{\mathtt{C}}, \hat{\mathtt{C}}'$, where $\hat{\mathtt{C}}, \hat{\mathtt{C}}'$ are consistent with the epipolar geometry. The quadric $\hat{\mathtt{Q}}$ provides a parametrization for this minimisation — its outlines, $\hat{\mathtt{C}}, \hat{\mathtt{C}}'$, in each image are generated via the camera projection matrices $\mathtt{P}, \mathtt{P}'$ (as in equation (5)); and outlines computed in this manner are consistent with the epipolar geometry.

Starting from the corrected conics as described above, one member of the quadric pencil is selected by choosing a 3D plane tangent to the quadric. The choice does not affect the reprojected outlines. The reprojection errors are are minimised by varying the remaining eight quadric parameters. The minimisation is carried out using the Levenberg-Marquardt algorithm.

3.5 Two view examples

Figure 1 shows an example of reconstruction from points correspondences alone—104 correspondences are used. The accuracy is demonstrated by the computed conic outlines and the quadric induced image transfer. Table 1 and section 7 gives a quantitative comparison of point and outline methods for these images.

Figure 4 shows an example of reconstruction from outlines and a single point correspondence. The chimneys in the power-station are close to hyperboloids.

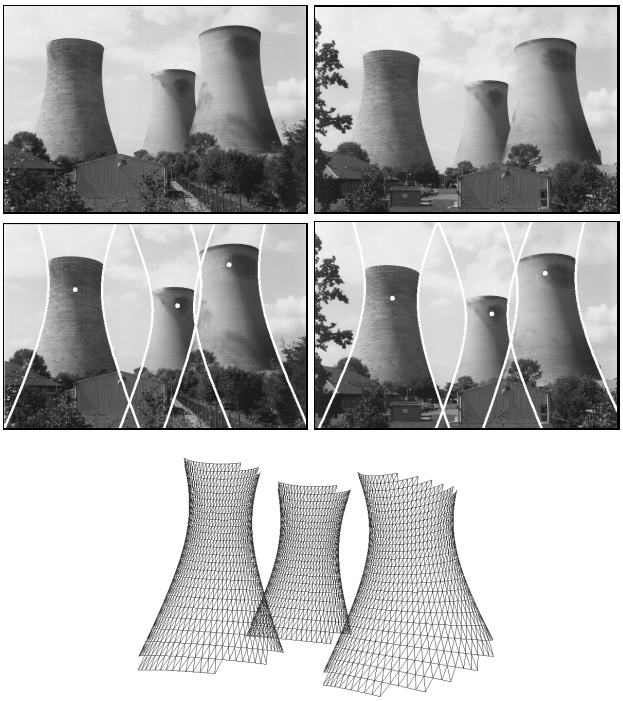


Figure 4: Two images of the hyperboloid chimneys at a power-station (*first* image set). The conic outlines and a single point correspondence on each quadric (*second* image set) allow the quadric to be reconstructed. The final image depicts the VRML model of the reconstructed quadrics.

4 Reconstruction from outlines in 3+ views

The image conic outline in a third applies a further set of constraints to the quadric. Following a similar argument to that outlined above, and assuming general positions of the camera: the first image conic imposes a full 5 constraints on the quadric; as the second image conic is constrained by the epipolar tangencies from the first image, it only imposes a further 3 constraints on the quadric; finally, the third image conic is constrained by the epipolar tangencies from both the first and second image, and hence only imposes a single constraint on the quadric.

Hence, there are a total of 9 constraints from three images, and the quadric is determined uniquely. Additional images improve the estimation accuracy.

A linear solution for the quadric is obtained by extending the matrix of equation (7) to the additional views. For 3 views $\mathtt{M}$ is a 18×13 matrix of rank 12. The (dual) quadric is obtained from the one dimensional null-space.

The method is implemented numerically in a similar manner to the two view case. Measured conics are first corrected to be consistent with the trifocal geometry constraints, and these corrected conics provide an initial solution. An optimal solution is then obtained using the same quadric parametrization as earlier, but with the cost function extended to the sum of squared perpendicular distance of the measured edgels from the projected conic outlines in all three (or more) views.

Figure 5 shows an example.

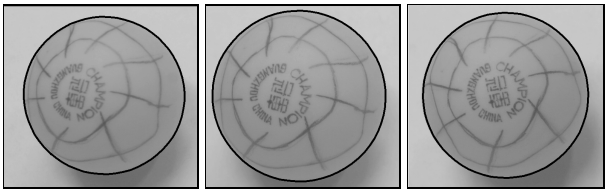


Figure 5: Three different views of a ping-pong ball. The quadric is computed using only the conic outlines (shown). The reconstruction is displayed in figure 7a.

5 Degeneracy Constraint

A degenerate quadric (e.g. a cone) is represented by a singular 4×4 matrix. The $\det \mathbf{Q} = 0$ constraint can be used to reduce the reconstruction ambiguity.

Reconstruction from points As explained in section 1, nine 3D points determine a quadric. However, it is possible to "replace" one of these points for a degenerate quadric by enforcing that the matrix of coefficients of the recovered quadric, $\mathbf{Q}$, is singular. In more detail 8 points define a one parameter of family of quadrics,

$$\mathbf{Q} = \mathbf{Q_A} + \lambda \mathbf{Q_B} \ ,$$

where $\mathbf{Q_A}, \mathbf{Q_B}$ are the quadrics corresponding to the vectors that span the null-space of $\mathtt{A}$. The constraint $\det \mathbf{Q} = 0$ produces a quartic polynomial in λ, and hence there are up to four real solutions for $\mathbf{Q}$.

Figure 6 shows an example.

6 Piecewise quadric reconstructions

Many objects can be accurately piecewise approximated by quadrics. This not only leads to an accurate model of the surface, but also to a compact representation.

The piecewise quadric representation is built from a piecewise conic fit to the image outlines as follows: the linked edgels composing the image outline are partitioned at curvature discontinuities; a conic is fitted to



Figure 6: Two images of a champagne bottle which locally is an example of a degenerate quadric surface (a cylinder). By imposing the degeneracy constraint, the reconstruction can be made from only 8 points (see figure 7b).

each segment using Bookstein's estimator [4]; the resulting conics are matched between the images and the epipolar tangency constraint imposed; the matched conics define the quadrics. In most cases, the number of conics to be matched is small. Where there is ambiguity, the selection is on a measure of closeness to the epipolar tangency constraint.

Figure 8 shows two images of a champagne bottle. The outline has been segmented into four conics, and each matched pair is reconstructed as a quadric to provide the overall reconstruction.

7 Comparison of reconstruction constraints

Table 1 gives a quantitative comparison of the different constraints arising from imaged outlines, imaged surface points, and a combination of the two. The example used is that of the ping-pong ball in figures 1 and 5. There are up to 75 point correspondences, and outlines in up to 3 views.

Two image measures are defined: for Hartley-Sturm corrected correspondences the conic induced *transfer error* is the average of the distances between a measured point and its transferred counterpart; the conic *outline error* is defined as the average of the perpendicular distances of measured outline points from the (optimally) estimated conic.

The results show that increasing the number of points improves the quadric reconstruction. Reconstruction from points alone gives accurate transfer between the two images, but the conic outline errors are significant. This is because an accurate quadric is fitted in the areas covered by point correspondences, but in the areas where point correspondences are sparse or not as well localized (such as close to the conic outline in the image) the quadric is not as well constrained. A single outline constraint reduces the outline error significantly, whilst leaving the transfer errors almost unchanged. A second outline constraint reduces the outline error to less than 0.1 pixels, whilst reducing the transfer error further.



Figure 7: Once the quadric has been reconstructed, a VRML model can be made (as shown), and novel views created. Texture is mapped onto the quadric from one of the images. The projective skew has been removed for this visualization. (a) A ping-pong ball reconstructed from outlines only in 3 views (see figure 5). (b) A champagne bottle reconstructed using the degeneracy constraint and conic outlines in 2 views (see figure 6).

8 Conclusions

The theory developed in this paper has been significantly simplified by adopting a dual-space approach. This not only allows the complicated multiview geometry to be visualised [6], but also enables extensions to reconstructing any surface with a known dual from multiple views. The natural extension is to the recovery of cubic surfaces, but other surfaces such as surfaces of revolution can also be approached.

Acknowledgements

We are grateful for software provided by Andrew Fitzgibbon, and for financial support from EU Project VANGUARD, the University of Oxford and Balliol College, Oxford.

References

[1] Bascle B., Blake A., and Zisserman A. Motion deblurring and super-resolution from an image sequence. In *Proc. 4th European Conf. on Computer Vision, Cambridge*, pages 312–320. Springer-Verlag, 1996.

[2] Beardlsey P., Zisserman A., and Murray D. Sequential update of projective and affine structure from motion. *International Journal of Computer Vision*, 23(3):235–259, 1997.

[3] Beardsley P., Torr P., and Zisserman A. 3D model acquisition from extended image sequences. In *Proc. European Conference on Computer Vision*, LNCS 1064/1065, pages 683–695. Springer-Verlag, 1996.

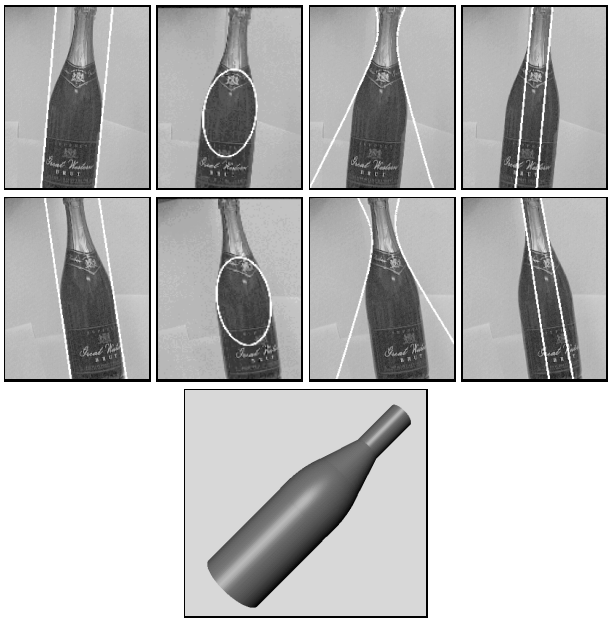


Figure 8: The outline of a champagne bottle segmented into 4 conics. *Upper images*: conics for the first view. *Middle images*: conics for the second view. *bottom image*: the piecewise quadric reconstruction of the champagne bottle.

# of outlines	# of points	Transfer error	Outline error
0	9	1.1×10^{0}	30.0
0	15	1.3×10^{-3}	12.0
0	50	4.1×10^{-4}	6.0
0	74	6.1×10^{-5}	3.0
1	4	2.6×10^{-1}	0.3
1	15	1.2×10^{-3}	0.3
1	50	3.2×10^{-4}	0.2
1	74	1.1×10^{-6}	0.1
2	(2)	2.3×10^{-3}	0.0
2	15	1.3×10^{-4}	0.0
2	50	7.3×10^{-5}	0.0
2	74	2.1×10^{-7}	0.0
3	0	8.3×10^{-3}	0.1
3	15	6.3×10^{-5}	0.0
3	50	2.1×10^{-5}	0.0
3	74	1.8×10^{-7}	0.0

Table 1: Errors in reconstruction for a varying number of point and outline constraints. Note: in the case of 2 conic outlines, and 1 image point, a further point is necessary to reduce the 2-fold ambiguity (see section 3), but this does not affect the accuracy of the fit.

[4] Bookstein F. Fitting conic sections to scattered data. *Computer Vision, Graphics and Image Processing*, 9:56–71, 1979.

[5] Canny J. A computational approach to edge detection. *IEEE Trans. on Patt. Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[6] Cross G. and Zisserman A. Quadric surface reconstruction from dual-space geometry. Technical report, Department of Engineering Science, University of Oxford, 1997.

[7] Hartley R. and Sturm P. Triangulation. In *Proc. Conference Computer Analysis of Images and Patterns*, Prague, Czech Republic, 1995.

[8] Irani M. and Peleg S. Improving resolution by image registration. *Graphical Models and Image Processing*, 53:231–239, 1991.

[9] Karl W., Verghese G., and Willsky A. Reconstructing ellipsoids from projections. *CVGIP: Graphical Models and Image Processing*, 56(2):124–139, March 1994.

[10] Luong Q.–T. and Viéville T. Canonic representations for the geometries of multiple projective views. In *Proc. 3rd European Conf. on Computer Vision, Stockholm*, pages 589–599, May 1994.

[11] Ma S.D. and Chen X. Quadric surface reconstruction from its occluding contours. In *Proc. International Conference on Pattern Recognition*, Israel, 1994.

[12] Ma S.D. and Li. L. Ellipsoid reconstruction from three perspective views. In *Proc. International Conference on Pattern Recognition*, Vienna, 1996.

[13] Ofek E., Shilat E., Rappoport A., and Werman M. Highlight and reflection-independent multiresolution textures from image sequences. Interal report, Institute of Computer Science, The Hebrew University of Jerusalem, Jerusalem, Israel, 1995.

[14] Quan L. Invariant of a pair of non-coplanar conics in space: Definition, geometric interpretation and computation. In *Proc. 5th Int'l Conf. on Computer Vision*, pages 926–931, 1995.

[15] Semple J. and Kneebone G. *Algebraic Projective Geometry*. Oxford University Press, 1979.

[16] Shashua A. and Toelg S. The quadric reference surface: Theory and applications. *International Journal of Computer Vision*, volume 23, issue 2, 1997.

[17] Torr P. and Zisserman A. Robust computation and parameterization of multiple view relations. In *International Conference on Computer Vision*, Bombay, India, 1998.

[18] Triggs B. Auto-calibration and the absolute quadric. In *Proc. Computer Vision and Pattern Recognition*, 1997.

[19] Zhang Z., Deriche R., Faugeras O., and Luong Q. A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry. *Artificial Intelligence*, 78:87–119, 1995.