

Uncalibrated 3D metric reconstruction and flattened texture acquisition from a single view of a surface of revolution

Carlo Colombo Alberto Del Bimbo Federico Pernici
Dipartimento di Sistemi e Informatica
Via Santa Marta 3, I-50139 Florence, Italy
{colombo,delbimbo,pernici}@dsi.unifi.it

Abstract

We describe a geometric approach for reconstructing 3D textured graphical models of surface of revolution (SOR) objects from a single uncalibrated view. Our approach is based on the fact that, for the object class of interest, the structure of the scene provides enough constraints for camera calibration even from a single view. Reconstruction (up to a scaling factor) of 3D shape is complemented with the extraction of flattened 2D texture, so as to support visual retrieval from 2D/3D cues and to generate realistic 3D visualization models. The approach developed is quite simple, yet accurate and robust; its applications range from the preservation, analysis and classification of cultural heritage, to advanced graphics and multimedia.

1. Introduction

When taking a photograph of a painted curved surface, its image appears distorted due to both surface geometry and perspective projection. This paper presents a method for reconstructing the image painted on a surface of revolution (SOR) and the metric 3D structure of the surface from a single uncalibrated view.

SOR are a special subclass of straight homogeneous generalized cylinders (SHGC), a surface class which has been extensively studied over the last two decades [1], [2]. The standard approach to the problem is that of trying to recover a volumetric description without actually modeling the camera; as a result, reconstruction of 3D structure cannot be said to be really metric. The same observation applies for previous approaches attempting at extracting the 2D scene depicted on the surface. For example, the work on paintings presented in [3] produces mosaic images that can have a different aspect ratio with respect to the original painting.

Our approach is mainly based on the application of projective geometry concepts to computer vision [4], [5], being specifically inspired by recent works on visual reconstruction from multiple images and camera-self calibration from scene constraints [6], [7], [8]. Previous approaches typically require two different images of the same SOR, or two distinct SOR objects inside a single image to calibrate the camera. Yet, our analysis shows that enforcing the constraint on imaged cross-section parallelism gives us enough information to calibrate the camera from just a single view of a SOR. We exploit this result in order to obtain a metric reconstruction (up to a scale factor) of a SOR from a single uncalibrated view.

Reconstruction of both 3D shape and flattened 2D texture is useful both to support visual retrieval from 2D/3D cues and to generate realistic 3D visualization models. The approach is quite simple, yet accurate and robust; its implementation features a simple graphical user interface allowing the users to interact with the system during model construction. Applications range from preservation and classification of cultural heritage, to advanced graphics and multimedia.

2. Algorithm and Implementation

2.1 Overview

Examples of SOR can be found among man-made objects of common use (pottery, cans, etc.) and in architectural design (e.g., parts of buildings). In the following, our attention will be focused on textured SOR objects, such as the decorated vase of Fig. 1, which are characterized both by their 3D shape and by their 2D pictorial content on their surface. Two lathe-crafted vases can have



Figure 1. Image of a decorated chinese vase.

exactly the same shape but different decorations (figurative vs geometric), but can also share the same decorations while being very different in shape (bell-like vs conic). The shape and pictorial elements are encoded in a very complex way into a single perspective image of a textured SOR; in the following, we will expose a geometric method to invert perspective projection, and infer both shape (up to a scale factor) and texture from a single *uncalibrated* image of the object. This is achieved in two steps. First, the projective properties of the SOR model are exploited so as to

reconstruct the actual object shape. This step also involves computing camera calibration from scene constraints as a way to get real metric information from images. Secondly, knowledge of shape will be used to perform visual texture acquisition.

2.2 Projective properties of SOR objects

Being a subclass of SHGC, SOR enjoy all of their properties [7]. A SOR can be parametrized as $\sigma(\theta, z) = (\rho(z) \cos \theta, \rho(z) \sin \theta, z)$, where z is the (straight) axis of revolution. In 3D space, all *parallels* (i.e., cross-sections with planes $z = \text{constant}$ orthogonal to the axis) are circles. The curves $\theta = \text{constant}$, called *meridians*, are obtained by cutting the surface with planes passing through the axis, and characterize the specific SOR shape through the *scaling function* $\rho(z)$. Parallels and meridians are locally mutually orthogonal in 3D space, but not in a 2D view (see Fig. 2). Typically observable curves in a SOR image are imaged parallels (which are always

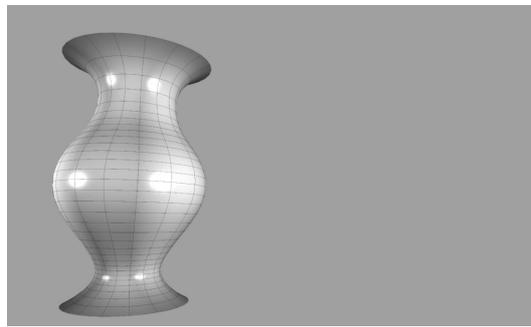


Figure 2. Parallels and meridians on a SOR.

ellipses, being the perspective images of circular curves) and *apparent contours* (see Fig. 3): the latter should not be confused with imaged meridians. In fact, while meridians are planar 3D

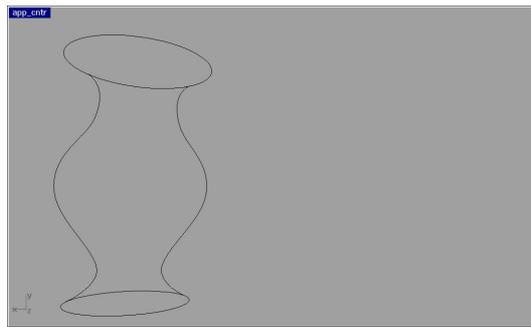


Figure 3. Ellipses and apparent contours for an imaged SOR.

curves, an apparent contour is the image of the (usually non planar) 3D curve of all the points at which the projection rays are tangent to the surface, referred to as *generating contour*. Fig. 4 remarks the difference between imaged meridians and apparent contours.

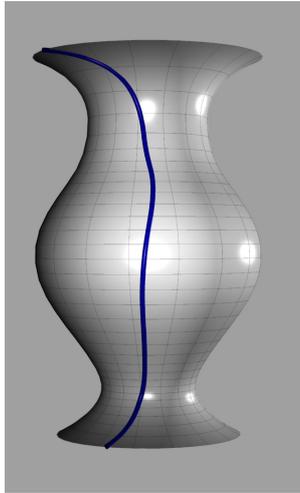


Figure 4. Generating contours and meridians generally differ. This makes apparent contours and imaged meridians also to be different.

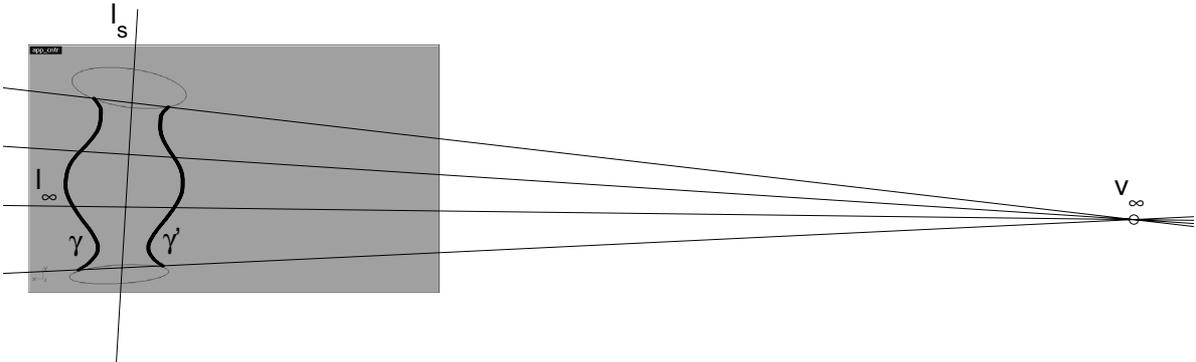


Figure 5. Visual symmetry under full perspective is governed by an harmonic homology.

Projective properties of SHGC and SOR can be conveniently illustrated by means of particular planar projective transformations called homologies [9], [4].

A plane projective transformation W is a planar homology if it has a line of fixed points \mathbf{l} (the axis) and a fixed point \mathbf{v} (the vertex) not on the line. This means that the associated matrix has two equal and one distinct eigenvalues, with their eigenspaces respectively of dimension two and one. The ratio of the distinct eigenvalues to the repeated one is the characteristic invariant μ . The projective transformation representing the homology can be parametrized in terms of $\mathbf{l}, \mathbf{v}, \mu$ as

$$W = I + (\mu - 1) \frac{\mathbf{v}\mathbf{l}^T}{\mathbf{v}^T\mathbf{l}} . \quad (1)$$

A planar homology thus has five degrees of freedom. A specialization of a planar homology is obtained when $\mu = -1$, in which case the homology is called harmonic, and has four degrees of freedom.

This said, we can state the four basic projective properties of a SOR:

1. Any two imaged parallels, say i and j , are related by a planar homology W_{ij} with the vertex \mathbf{v}_{ij} on the imaged axis of the SOR, the axis \mathbf{l}_∞ being the vanishing line of the planes orthogonal to it:

$$W_{ij} = I + (\mu_{ij} - 1) \frac{\mathbf{v}_{ij}\mathbf{l}_\infty^T}{\mathbf{v}_{ij}^T\mathbf{l}_\infty} \quad (2)$$

2. Imaged meridians are related by a harmonic homology H whose axis is the imaged SOR axis and whose vertex is on the vanishing line of the planes orthogonal to the imaged SOR axis. It holds

$$H = I - 2 \frac{\mathbf{v}_\infty\mathbf{l}_s^T}{\mathbf{v}_\infty^T\mathbf{l}_s} \quad (3)$$

3. The apparent contour is tangent to the imaged parallels at the point of contact.
4. The two sides of the apparent contour of a SOR are related by an harmonic homology, whose axis is the imaged SOR axis (as in property 2 above).

Of the properties above, the first three apply to the more general class of SHGC objects, while the fourth is characteristic of SOR objects, being an extension of the second property. Specifically, the fourth property illustrates in which terms the usual concept of bilateral symmetry has to be revised in the presence of projective deformations (see also Fig. 5).

2.3 3D reconstruction

In the case of a SOR, the problem of 3D reconstruction is equivalent to that of reconstructing the scaling function, i.e. the shape of any meridian. This can be achieved in three main steps: (1) evaluation of the two homologies, namely W_{ij} and H , referred to in the first and fourth properties above, and determination of an imaged meridian; (2) camera calibration by exploitation of SOR scene constraints; (3) metric rectification of the imaged meridian.

The input data for the algorithm are assumed to be two elliptical imaged parallels, referred to as C_1 and C_2 and represented as usual as 3×3 symmetric matrices, and one (side of the)

apparent contour, γ (refer again to Fig. 3). The latter input is actually used only in the first step, to estimate the imaged meridian; also, only one of the two input ellipses is strictly needed to accomplish our goal [10]: we have elected to work with two distinct ellipses only for the sake of algorithmic robustness.

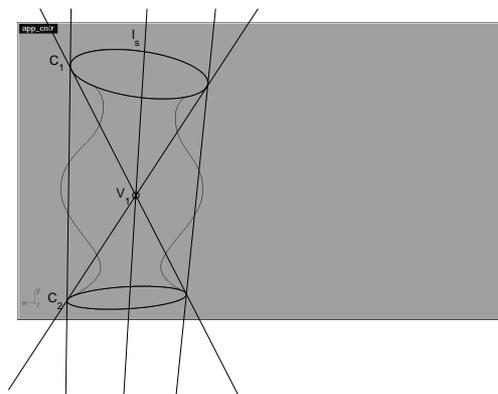


Figure 6. Computing the axis of symmetry l_s of the harmonic homology.

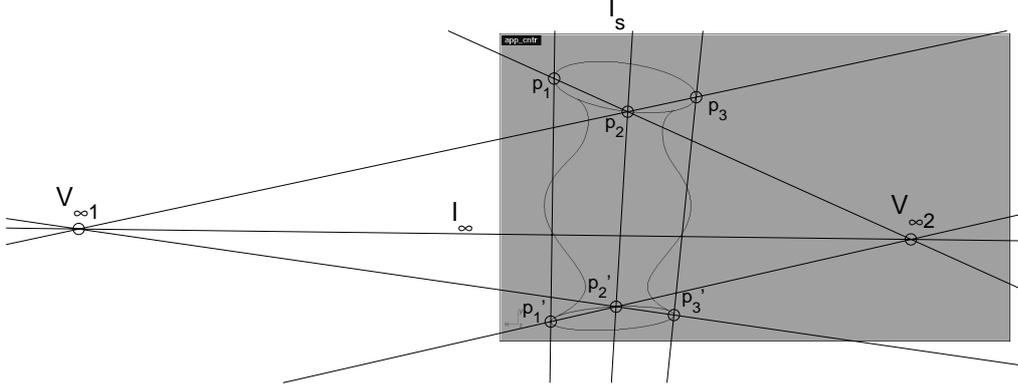


Figure 7. Computing the vanishing line l_∞ of the planar homology.

Homology evaluation and imaged meridian estimation. The four degrees of freedom of the harmonic homology H (see fourth property above) are encoded in the imaged axis of symmetry \mathbf{l}_s (2 d.o.f.) and the vanishing point \mathbf{v}_∞ (2 d.o.f.).

To compute the imaged axis of symmetry, we notice that, although property 2 above is stated in terms of imaged meridians, it can be restated in an operational way in terms of imaged parallels, as each imaged parallel is transformed by H onto itself. This allows us using the two ellipses C_1 and C_2 and evaluating \mathbf{l}_s as the line through the points of intersection of the bitangents to the curves (see Fig. 6). That done, the vertex is easily computed as $\mathbf{v}_\infty = C_1^{-1}\mathbf{l}_s$, thanks to the pole-polar relationship with the axis of symmetry through any of the ellipses.

Another important quantity to be computed is the line at infinity of the planar homology W_{ij} , \mathbf{v}_∞ (2 d.o.f.), which is actually independent of i and j , i.e. of the choice of the ellipse pair. The two ellipses C_1 and C_2 can thus be used as shown in the geometric construction of Fig. 7, involving three corresponding point pairs, obtained by the contact of the two ellipses with the two bitangents and the axis of symmetry (three different solutions exist in this last case, corresponding to three different views of the SOR: the user must indicate to the algorithm the right one).

We are now in the position to use the ellipse C_1 and the apparent contour γ and estimate the imaged meridian corresponding to a reference direction α , measured in the image starting from the imaged symmetry axis \mathbf{l}_s . Let $\mathbf{q}_\alpha \in C_1$ be the point of C_1 corresponding to the direction α . The algorithm, to be iterated for any $\mathbf{p}_k \in \gamma$, is the following:

1. use the tangency property 3 above to find the vanishing point \mathbf{v}_k as the intersection of the tangent to γ at \mathbf{p}_k and \mathbf{l}_∞ ;
2. compute the tangent to C_1 passing through \mathbf{v}_k : the tangency point \mathbf{q}_k corresponds to \mathbf{p}_k via the (still unknown) planar homology W_{1k} ;
3. draw the line through \mathbf{p}_k and \mathbf{q}_k : it intersects (in any order) the horizon line \mathbf{l}_∞ in \mathbf{r}_k , and the imaged axis of symmetry \mathbf{l}_s in \mathbf{v}_{1k} ;

4. compute the characteristic invariant μ_{1k} as the cross-ratio of the four points \mathbf{p}_k , \mathbf{q}_k , \mathbf{r}_k and \mathbf{v}_{1k} . This concludes the evaluation of W_{1k} ;
5. map $\mathbf{q}_\alpha \in C_1$ through W_{1k} to obtain the desired imaged meridian point, $\mathbf{p}_\alpha \in C_k$.

Fig. 8 shows several image meridians computed for and superimposed to the chinese vase of Fig. 1.

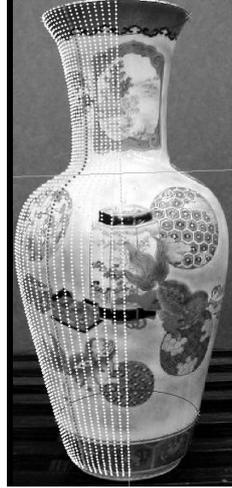


Figure 8. Imaged meridians for the chinese vase of Fig. 1.

Metric rectification of the imaged meridian. Let us assume for the moment to know already the internal camera parameters, i.e. to know the entries of the calibration matrix

$$\mathbf{K} = \begin{pmatrix} f & s & u_0 \\ 0 & rf & v_0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (4)$$

where f is the focal length, r is the pixel aspect ratio, u_0 and v_0 are the coordinates of the principal point and s depends on the skew angle between the coordinate image axes. The image of the absolute conic (IAC), an imaginary point conic usually referred to as ω , is a geometric tool often used in the place of \mathbf{K} to carry information about internal camera parameters: it holds $\omega = \mathbf{K}^{-T}\mathbf{K}^{-1}$. Given the IAC, it is well known that any plane for which the vanishing line is known can be rectified [11]. The geometric entities involved in planar rectification are shown in Fig. 9, where it is assumed that $r = 1$ (square pixels) and $s = 0$ (no skew). Under this assumption, the IAC is reduced to a circle, as it depends on only three internal parameters. Being the IAC a complex-valued curve, Fig. 9 shows in its place a real-valued conic (in our case, a circle), called calibrating conic, which is a geometric device for IAC visualization [4]. The figure also shows two vanishing lines: \mathbf{l}_∞ represents any plane perpendicular to the SOR axis, while $\mathbf{l}_{\infty\rho}$ represents any plane parallel to the meridian we want to rectify. Let us associate an image direction β to this meridian: to evaluate $\mathbf{l}_{\infty\rho}$, all we need is to join the vanishing point $\mathbf{v}_{\infty 1}$ relative to the direction

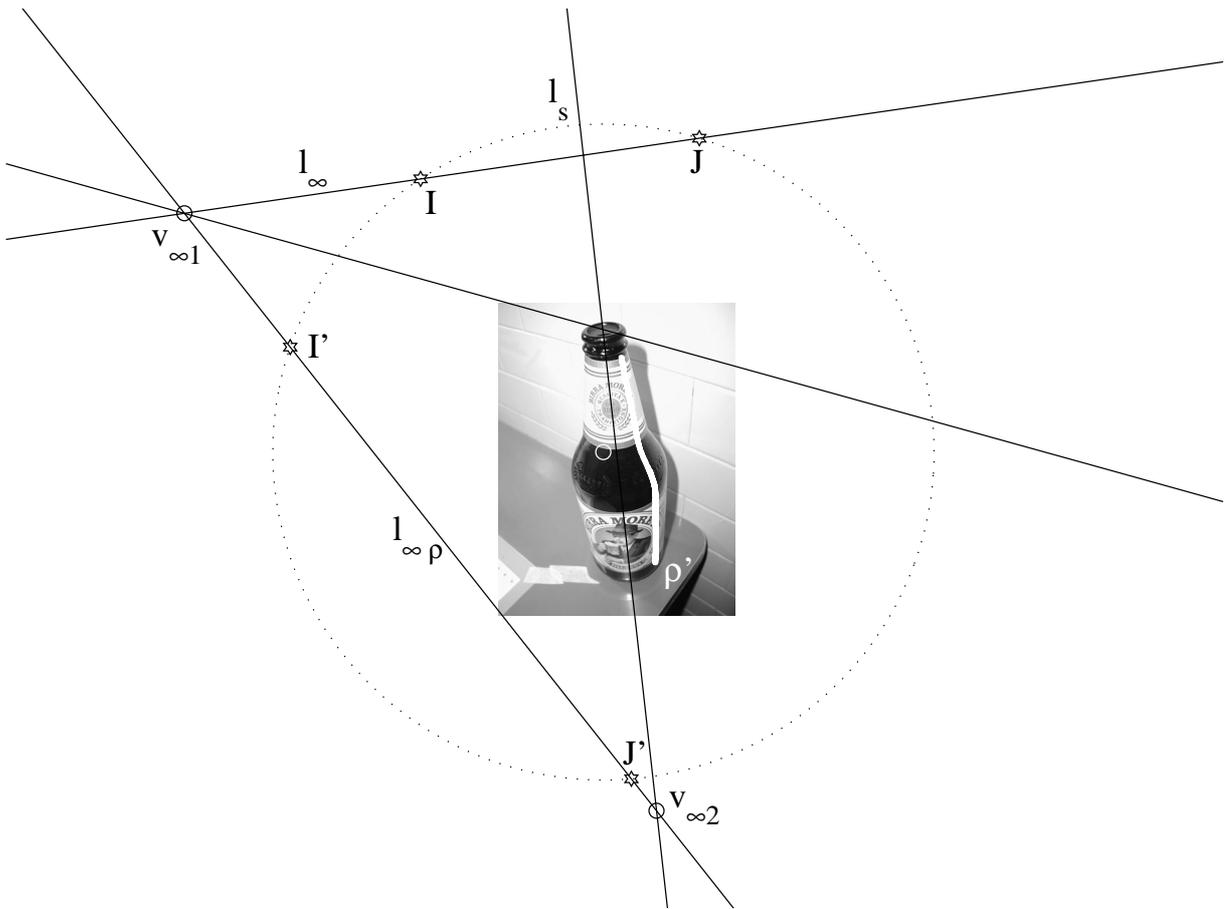


Figure 9. The geometry of 3D reconstruction and camera calibration.

β with the vanishing point $\mathbf{v}_{\infty 2} = \omega^* \mathbf{l}_{\infty}$ relative to the direction orthogonal to all parallels (ω^* is the adjoint matrix of ω , also known as Dual IAC, or DIAC).

Intersecting the vanishing line for a given plane with the IAC yields two complex conjugate points (in the figure, \mathbf{I} and \mathbf{J} are related to \mathbf{l}_{∞} , while \mathbf{I}' and \mathbf{J}' are related to $\mathbf{l}_{\infty \rho}$), which are known as *circular points*: this requires solving a quadratic equation [9]. Once $\mathbf{l}_{\infty \rho} = (l_1, l_2, 1)$ and $\mathbf{I}' = (a, 1, b)$ are computed, the rectifying homography for the meridian relative to the direction β can be calculated as [12]

$$\mathbf{H}_{\beta} = \begin{pmatrix} 1/d & -c/d & 0 \\ 0 & 1 & 0 \\ l_1 & l_2 & 1 \end{pmatrix}, \quad (5)$$

where $c = \text{Re}(a)$, $d = -\text{Im}(a)$. The rectifying homography \mathbf{H}_{β} allows us to unwarp from the projective distortion the imaged meridian of Fig. 9 and get finally the scaling function $\rho(z)$ of Fig. 10.



Figure 10. Recovered scaling function for the object of Fig. 9.

Camera calibration from SOR scene constraints. This paragraph shows how to calibrate the camera from SOR scene constraints, thus being able to use the IAC for metric rectification. Assuming that $r = 1$ and $s = 0$, then ω is a circle, depending on three parameters (refer again to Fig. 9). Hence, a minimum of three constraints must be set in order to estimate the IAC. In our method four constraints are set, and the resulting over-determined system is solved by least squares. The first two constraints derive from computing explicitly the circular points \mathbf{I} and \mathbf{J} as the (complex) intersection of one of the imaged parallels (say, C_1) with \mathbf{l}_{∞} : in fact, this is a basic property of circular points. The other two constraints come from the pole-polar relationship of \mathbf{l}_s and \mathbf{v}_{∞} through ω , i.e. $\mathbf{l}_s = \omega \mathbf{v}_{\infty}$.

2.4 Flattened texture acquisition

This section shows how to exploit the knowledge of the scaling function so as to perform texture acquisition in the form of a flat (planar) image. Flat images are required to extend the object representation and allow using traditional image database technology for the retrieval of textured

3D objects. Another objective is to build texture maps which are compliant with commercial 3D modeling-rendering software standards.

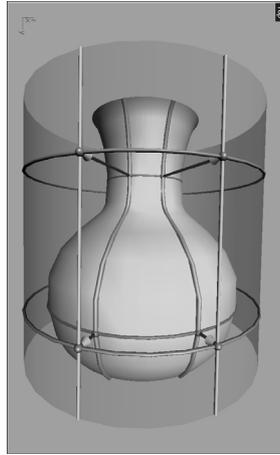


Figure 11. The normal projection map used for metric texture flattening.

We have elected to use one of the most common image mapping techniques, the *normal cylindrical projection*, as it fits well with the kind of objects we are dealing with. Cylindrical maps are constructed by aligning a sampling cylinder so that it is coaxial with the SOR (see Fig. 11). In such a way, parallels and meridians transform as a rectangular grid, meridians being equally spaced. The 3D reconstruction information is used, of course, in order to move metrically over the imaged meridians and parallels and map them to the texture image space. In the case of SOR, the normal cylindrical projection map is readily obtained from the natural surface parametrization, where θ and z are the texture coordinates.



Figure 12. A Pepsi can.

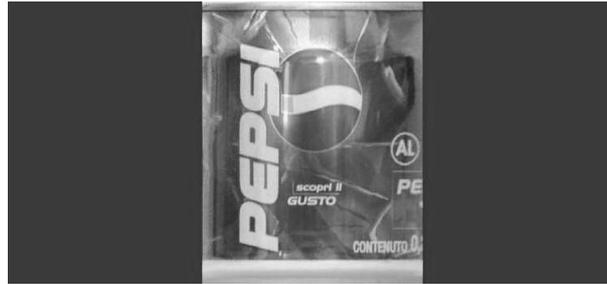


Figure 13. The flattened texture acquired from Fig. 12.

Fig. 12 shows the image of textured SOR (a tin can); the flattened texture is given in Fig. 13. Notice that, since in this case the SOR is a cylindrical (thus developable) surface, the flattened texture preserves the local geometry of the original 3D object, thus allowing for example to reproduce correctly as circular the 'AL' mark, even though such mark is highly deformed and nearly invisible in the original photo.

3. Tests and Application Examples

The algorithms were developed in C and run on a standard PC with Windows. A simple graphical user interface allows the user to load an image, select manually two ellipses and a portion of apparent contour, and run the program. The output of the program consists of a (1) high-resolution description of the scaling function, (2) a color image containing the flattened texture, and (3) a VRML model for immediate visualization of the reconstruction results.

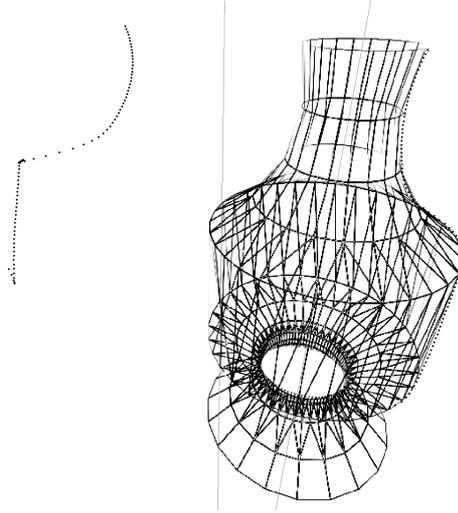


Figure 14. A synthetic model and its reconstructed scaling function.

In order to test the effectiveness and robustness of the above algorithm, several tests have been performed both on synthetic and on real images. Fig. 14 shows a synthetic model and, at its left, the reconstructed scaling function. The use of a synthetic ground truth allows us to estimate model accuracy, which is quite high. Actually, the departure of the reconstructed model from the true one is negligible, provided that user manual input is reasonably accurate.

The method has several applications in the field of cultural heritage, for instance for the classification and retrieval of archaeological vases, or for the visual reconstruction of frescoes. Fig. 15 reports the results of the geometric reconstruction of a painted circular vault. The pictorial content of the fresco can be extracted from the photograph, and developed on a planar image without loss of information.

The method can also be used to detect anomalies in images. For instance, hand-crafted vases are typically not perfect SOR objects, and this can be revealed by the reconstruction engine. Also, since an intermediate step of the method is to compute calibration parameters, is it possible to know whether the image being analyzed has been cropped from a larger image, by checking if the principal point lies inside or outside the image. For example, this is the case for Fig. 1, whose 3D model is reported in Fig. 16 (texture portions invisible in the original image are shown in black).

Fig. 17 shows two VRML views of the reconstructed model for the SOR object shown in Fig. 12. Once extracted, a visual texture can be used in combination with other 3D models, to yield curious results. An example of this application is reported in Fig. 18, showing a hybrid model obtained by combining the visual properties of Fig. 12 and Fig. 1.

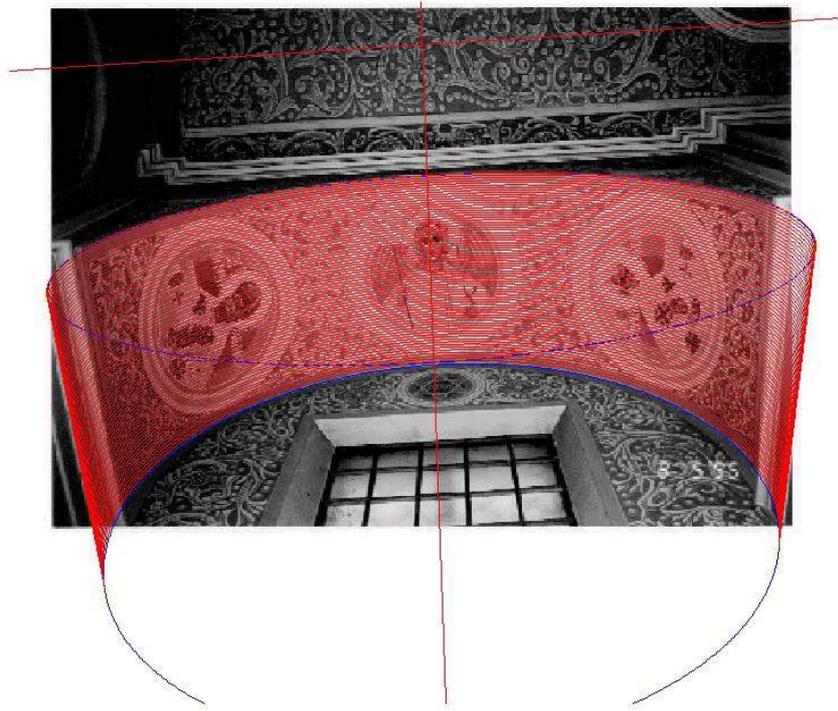


Figure 15. Reconstruction of the geometry of a circular vault.



Figure 16. Two synthetic views of the chinese vase reconstructed from Fig. 1.



Figure 17. Two synthetic views of the Pepsi can of Fig. 12.

4. Future Work

The current system is being expanded in two directions. On the one hand, it is being incorporated into a 2D/3D visual database engine, allowing to retrieve images and objects by 3D models and/or pictorial content. On the other hand, the algorithms are being improved in order to deal with more general object shapes, to run in a completely automatic way, and to be able to process multiple images of the same object, taken from different views.



Figure 18. The texture extracted from Fig. 12 is used to decorate the 3D model reconstructed from Fig. 1.

References

- [1] Jean Ponce, David Chelberg, and Wallace B. Mann, “Invariant properties of straight homogeneous generalized cylinders and their contours,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-11, no. 9, pp. 951–966, 1989.
- [2] M. Zerroug and R. Nevatia, “Part-based 3D descriptions of complex objects from a single image,” *PAMI*, vol. 21, no. 9, pp. 835–848, September 1999.
- [3] Adrian G. Bors, William Puech, Ioannis Pitas, and Jean-Marc Chassery, “Mosaicing of flattened images from straight homogeneous generalized cylinders,” in *Computer Analysis of Images and Patterns*, 1997, pp. 122–129.
- [4] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, 2000.
- [5] Olivier Faugeras and Quan-Tuan Luong, *The Geometry of Multiple Images*, MIT press, 2001.
- [6] D. Liebowitz, *Camera Calibration and Reconstruction of Geometry from Images*, Ph.D. thesis, University of Oxford, UK, Dept. of Engineering Science, June 2001.
- [7] Samer M. Abdallah, *Object Recognition via Invariance*, Ph.D. thesis, The University of Sydney, Australia, June 2000.
- [8] P.R.S. Mendonça, *Multiview Geometry: Profiles and Self-Calibration*, Ph.D. thesis, University of Cambridge, UK, May 2001.
- [9] J. Semple and G. Kneebone, *Algebraic projective geometry*, Oxford University Press, 1952.
- [10] F. Pernici, “Ricostruzione da singola immagine della struttura 3D e della tessitura superficiale di solidi di rivoluzione,” Master’s thesis, University of Florence, Italy, April 2002 (In Italian).
- [11] D. Liebowitz, A. Criminisi, and A. Zisserman, “Creating architectural models from images,” in *Proc. EuroGraphics*, September 1999, vol. 18, pp. 39–50.
- [12] D. Liebowitz and A. Zisserman, “Combining scene and auto-calibration constraints,” in *Proc. 7th International Conference on Computer Vision, Kerkyra, Greece*, September 1999.