

Shape-from-Shading under Near Point Lighting and Partial views for Orthopedic Endoscopy

Chenyu Wu[†] Srinivasa G. Narasimhan[†] Branislav Jaramaz^{† ‡}
[†] Carnegie Mellon University [‡] Institute of Computer Assisted Surgery
{cywu, srinivas}@cs.cmu.edu branko@icaos.org

Abstract

Bone reconstruction using endoscopy is important for computer aided minimally invasive orthopedic surgery. During surgery an endoscope consisting of a camera and one or more light sources is inserted through a small incision into the body and the acquired images are analyzed. Since bone surface is featureless, shading is the primary cue for shape perception. However, due to the small field of view of the endoscope, only a small part of the bone and its occluding contour are visible in any single image. Therefore even human perception of bone shape from such images can be hard.

We present a novel technique to reconstruct the surface of the bone by applying shape-from-shading to a sequence of endoscopic images, with partial boundary in each image. We first perform geometric and the photometric calibration for the endoscope. We then extend the classical shape-from-shading algorithm to include a near point light source that is not optically co-located with the camera. By tracking the endoscope we are able to align partial shapes obtained from different images in the global (world) coordinates. An ICP algorithm is then used to improve the matching, resulting in a complete occluding boundary of the bone. Finally, a complete and consistent shape is obtained by simultaneously re-growing surface normals and depths in all views. We demonstrate the accuracy of our technique using simulations and experiments with artificial bones.

1. Introduction

One of the main goals of orthopedic surgery is to enable minimally invasive procedures. As a key tool, endoscopy is attracting increasing attention for its potential role in computer aided surgery. Besides visualizing the interior of the anatomy, its role can be significantly enhanced by tracking the endoscope in conjunction with surgical navigation systems.

Fig 1 shows an endoscope illuminating and observing an

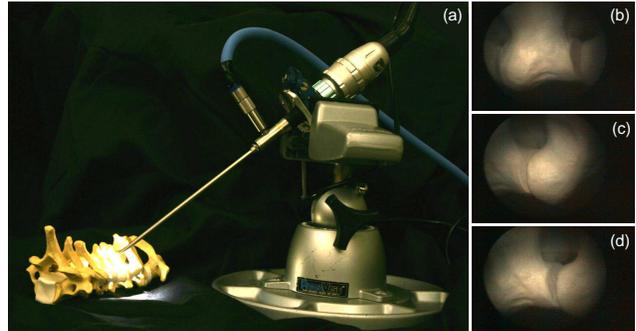


Figure 1. Illustration of an endoscope (a) and endoscopic images of an artificial spine (b,c,d). Small field of view, distortion, lack of texture, partial boundaries, and shading caused by near-field lighting make it difficult to perceive the shape of the object in the images.

artificial spine. Since the field of view of the endoscope is small and the bone is usually a few millimeters away, we only observe a small part of the bone. The endoscope uses one or more point light sources placed at its tip for illumination. Since the sources, camera and the scene are so close to each other, the endoscopic images are very different from the images we experience of common scenes under distant lighting such as from the sun or the sky. As a result, it can be difficult even for a skilled surgeon to infer bone shape from a single endoscopic image. Better visualization is achieved by overlaying endoscopic images with 3D surfaces obtained from CT scans [2, 4]. However, this still requires us to solve a complex registration problem between CT and endoscopic images during surgery. Thus, there is an immediate need for explicit computer reconstruction of bone shapes from endoscopic images.

Since bone surfaces have few identifiable features, surface shading is the primary cue for shape. Shape-from-shading has a long and rich history in computer vision and biological perception, with most of the work focusing on distant lighting [8, 12, 29]. Recently, several works have extended the traditional techniques to perspective projec-

tion [19, 23, 24] and also a near-point light source [3, 20]. Most relevant to this work, shape-from-shading under a near point light source and perspective projection for an endoscope has been proposed [3, 5]. But all these works assume the light source and camera center are co-located. Note that this assumption is inaccurate in our setting since the scene is 5-10 mm away from the endoscope and the source and camera separation distance (3.5mm) is of the same magnitude as the distance to the scene. Prados et al. [18] proposed a method based on the notion of Crandall-Lions viscosity solution, to obtain shape without complete image boundary. However, in our images where only a very small part of the boundary is visible, their technique is not enough to accurately recover the entire bone. Photometric stereo can be used to estimate surface orientation [26] but it is hard to achieve with a conventional endoscope since camera and sources are designed to be switched on or off at the same time and almost all endoscopes have at most 2 sources. Finally, structured light has also been used to recover bone shape [10] but requires specialized setups that are expensive and cumbersome to use in the operating room.

Due to the small field of view of the endoscope and the lack of a complete occluding contour, only a partial (and approximate) shape from a single image can be obtained. By capturing image sequences as the endoscope is moved, it is possible to cover a larger part of the shape that can be more easily perceived [21]. Once again, there is a long and rich history of structure-from-motion in vision (both calibrated and uncalibrated) [17, 16, 7]. In the context of anatomical reconstruction, the heart coronary [14] and tissue phantom [22] have been reconstructed using motion cues. However, note that it is difficult to achieve this for textureless (or featureless) bones.

So, neither shape-from-shading nor shape-from-motion can individually solve the problem of bone reconstruction from endoscopic images. The key idea in this work is to utilize both strengths to develop a global shape-from-shading approach using multiple partial views. We have two contributions in this paper. First, for a single image, we formulate shape-from-shading under a near-point light source and perspective projection, given only a partial object boundary, when the source and camera are not optically co-located. By taking into account the intensity fall-off from the source, we extend previous methods to estimate not only surface normals but depths as well. Second, for a sequence of images, we track the endoscope during image acquisition to be able to transform the reconstructed shape from a local view to a global (world) coordinate frame. The bone is immobilized so that the relative position of the endoscope to the bone is known at all times. Due to the error in the tracking system, the calibration process and the reconstruction caused by the partial boundary, the individual shapes from each image are not well aligned. We use ICP to match them

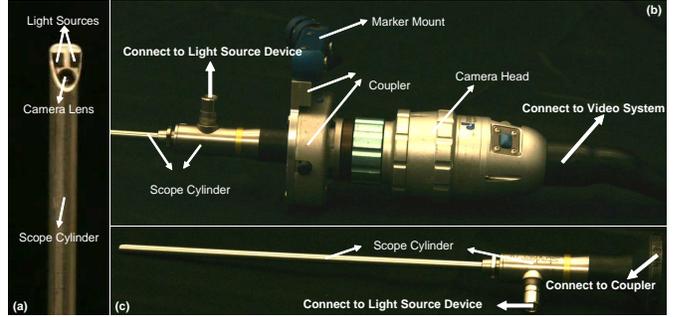


Figure 2. An oblique endoscope consists of a scope cylinder with a lens and two point light sources at the tip (the tip has a tilt from the scope cylinder) and a camera head that captures video. The scope cylinder is connected flexibly to the camera head via a coupler. The scope cylinder or the camera head can be rotated together or separately.

further to obtain the entire occluding boundary in world coordinates. We then restart the shape-from-shading for all images simultaneously and use boundary constraints in world coordinates in each iteration. This process of growing shape locally and updating constraints globally is iterated until convergence. As a clarification, note that our global shape-from-shading framework is not really merging shape-from-shading and shape-from-motion in the classical sense of the terms.

Before we can apply the above algorithm, the endoscope must be carefully calibrated. This is an important step that effects the accuracy of the final shape. Thus, we first present a calibration technique for obtaining both geometric and photometric parameters of the endoscope. While most of the technical details of calibration are presented in our technical report [27], we focus here on the issues that are specific to the endoscope that are not usually valid for regular cameras. We believe that our techniques can significantly enhance the capabilities of endoscopy in minimally invasive orthopaedic surgery.

2. Endoscope Geometry and Photometry

The orthopedic endoscope has a single camera and one or more point light sources equipped at the tip of the scope. For this work, we use the Stryker 344-71 arthroscope Vista (70 degree, 4mm), an oblique endoscope with two point light sources. The device is shown in Fig. 2. In this section, we briefly describe the geometric and photometric calibration required for this endoscope.

2.1. Geometric Calibration

Unlike regular cameras with forward viewing, most current endoscopes are designed for oblique viewing with a tilt from the scope axis at the tip (see Fig. 2 (a)). The view-

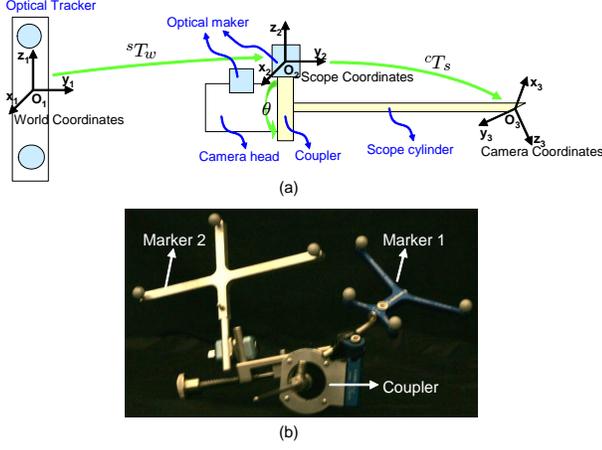


Figure 3. (a) The geometric model of endoscope based on a tracking system. (b) Two optical markers are attached to the scope cylinder and camera head and tracked separately.

ing direction can be changed by simply rotating the scope cylinder without having to move it. To increase flexibility of imaging in terms of focus and resolution, the lens system (in the cylinder) and the CCD plane (in the camera head) can be rotated with respect to one another (see Fig. 2 and 3 (a)). While they are desirable, these properties make the geometric calibration harder than in regular cameras.

Yamaguchi et al. [28] calibrate an oblique scope by extending Tsai’s model [25, 30] by attaching an optical marker to the camera head. However, since there can be a rotation between the head and the cylinder, the estimation of this rotation angle remains hard. Instead, we attach the marker on the scope cylinder and the calibration turns out to be simpler. The geometric model is illustrated in Fig. 3 and can be formulated as:

$$\begin{aligned} \lambda p'_i &= A \cdot {}^cT_m \cdot {}^mT_w \cdot P_w \\ p_i &= R(\theta) \cdot (p'_i - cc) + cc \end{aligned} \quad (1)$$

where, P_w is a 3D point in the world coordinates, p'_i is the corresponding 2D image pixel without rotation, p_i is the image pixel with rotation θ . mT_w is a rigid transformation from world coordinates to the marker coordinates, cT_m is a rigid transformation from the marker to the camera coordinates and cc is the camera center. Camera intrinsic matrix A and cT_m can be calibrated by using Zhang’s method [30] and mT_w can be obtained directly from the tracking system. Without a rotary encoder, the rotation angle can be estimated by using an extra marker (marker2 attached to the head as shown in Fig. 3).

2.2. Photometric Calibration

In addition to computing the radiometric response function of the camera in the endoscope, we must also simul-

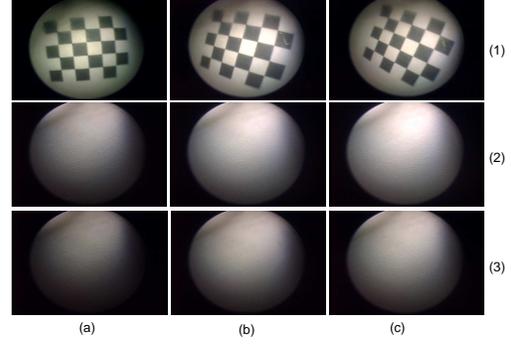


Figure 4. Images used for geometric calibration (1) and photometric calibration (2-3). Row 2-3 show different light intensities and Column a-c show different albedos on a color-checker.

taneously calibrate the directional/spatial intensity distribution of the near field sources (the camera and the sources are “on/off” at the same time). We derive a method similar in spirit to [13] to estimate the camera response function and spatial distribution of light intensities assuming the Lambertian model for reflectance.

The endoscope projection with two near point light sources is shown in Fig. 5. The irradiance received by the camera is written as:

$$R(x, y, z, f, g) = E_0 \rho \left(\frac{\mathbf{n}(f, g) \cdot \mathbf{l}_1(x, y, z)}{r_1(x, y, z)^2} + \frac{\mathbf{n}(f, g) \cdot \mathbf{l}_2(x, y, z)}{r_2(x, y, z)^2} \right) \quad (2)$$

where, E_0 is the intensity of sources s_1 and s_2 , ρ is the surface albedo and \mathbf{n} is the surface normal. \mathbf{l}_1 and \mathbf{l}_2 are two light rays incident at the surface and r_1 and r_2 are the distances from each light source to the surface. (f, g) is the stereographic transformation of the normal. (x, y, z) indicates the 3D location of the scene point P in Fig. 5. This projection model can be extended for multiple point light sources such as a ring light source. The detailed derivation is given in our technical report [27]. The image irradiance E is related to the image intensity v using the camera response function $H(\cdot)$:

$$E(x, y) = \frac{H^{-1}(v(x, y))}{M(x, y)} \quad (3)$$

where, $M(x, y)$ represents the anisotropy of the source intensity. The two sources are identical and are oriented in the same way so we assume their intensity distributions $M(x, y)$ are the same. From Equation 2 and 3 and image irradiance equation we have:

$$\begin{aligned} H^{-1}[v(x, y)] &= \rho \cdot E_0 \cdot \tilde{M}(x, y) \\ \tilde{M}(x, y) &= M(x, y) \cdot \left(\frac{\mathbf{n} \cdot \mathbf{l}_1}{r_1^2} + \frac{\mathbf{n} \cdot \mathbf{l}_2}{r_2^2} \right) \end{aligned} \quad (4)$$

For calibration, we use a Macbeth color chart with known albedo for each patch. We capture a set of images

by varying the source intensity for each patch (the exposure time of the camera cannot be controlled in our endoscope). We apply log to both sides of Equation 4 to obtain a linear system:

$$h[v_i^j(x, y)] = p_i + g_j + \tilde{m}(x, y) \quad (5)$$

where, i indicates the different albedos and j indexes the light intensities. $h[v_i^j(x, y)] = \log\{H^{-1}[v_i^j(x, y)]\}$, $p_i = \log(\rho_i)$, $g_j = \log(E_{0j})$ and $\tilde{m}(x, y) = \log[M(x, y)]$. The unknowns ($h(\cdot)$, g_j , $\tilde{m}(x, y)$) can be obtained by solving this linear system of equations. The term $M(x, y)$ is then estimated by physically measuring the distance to the chart from the scope tip.

3. Shape-from-Shading under Near Point Light Sources and Perspective Projection

By modeling the endoscope imaging system using perspective projection with near point light sources, shape-from-shading can be achieved by minimizing the error between the image brightness $E(x, y)$ and the reflectance map R . R is computed based on the model illustrated in Fig. 5 and the derivation is given in our technical report [27]. The important difference compared to the classical shape-from-shading [9] is that we take into account the intensity fall-off from a divergent source, so R is not only a function of normal but also a function of distance z (Prados and Faugeras [20] consider the fall-off from a distant source). The error can be computed as:

$$\mathbf{e}_i(z, f, g) = \int \int_{image} [E(x, y) - R(x, y, z, f, g)]^2 dx dy \quad (6)$$

By adding smoothness constraints for z , f and g as:

$$\mathbf{e}_s(z, f, g) = \int \int_{image} [(z_x^2 + z_y^2) + (f_x^2 + f_y^2) + (g_x^2 + g_y^2)] dx dy \quad (7)$$

we thus obtain a solution to the depth and normal:

$$[z^*, f^*, g^*] = \arg \min_{z, f, g} (\lambda \mathbf{e}_i + (1 - \lambda) \mathbf{e}_s) \quad (8)$$

There are several methods for solving image-irradiance equation such as using first order PDE [8], level-set methods [3, 5, 15], and viscosity solution [11, 19]. While we still use first order PDE for solving the image irradiance equation, we simultaneously propagate depth z during iterations. We discretize Equation 8 and obtain:

$$\mathbf{e}(z, f, g) = \sum_i \sum_j (\lambda \mathbf{e}_{i,j} + (1 - \lambda) \mathbf{e}_{s,i,j}) \quad (9)$$

where $\mathbf{e}_{i,k,l}$ represents the irradiance error at pixel (k, l) and $\mathbf{e}_{s,i,j}$ defines the smooth constraints at pixel (i, j) . λ is a Lagrange multiplier. The solution to Equation 9 is to find $\{f(k, l)\}$, $\{g(k, l)\}$ and $\{z(k, l)\}$ that minimize \mathbf{e} :

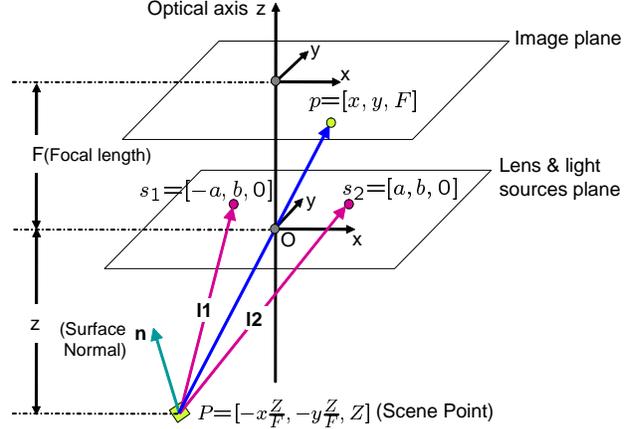


Figure 5. Perspective projection model for endoscope imaging system with two near point light sources: O is the camera projection center. s_1 and s_2 are two light sources. We assume the plane consisting of O , s_1 and s_2 is parallel to the image plane. The coordinate system is centered at O and Z is parallel to the optical axis, and pointing toward the image plane. X and Y are parallel to the image plane. F is the focal length. a and b are two parameters related to the position of the light sources. Given a scene point P , the corresponding image pixel is p . Assuming a Lambertian surface, the surface illumination therefore depends on the surface albedo, light source intensity and fall off, and the angle between the normal and light rays.

$$\frac{\partial \mathbf{e}}{\partial f_{k,l}} = 0, \frac{\partial \mathbf{e}}{\partial g_{k,l}} = 0, \frac{\partial \mathbf{e}}{\partial z_{k,l}} = 0 \quad (10)$$

Combining Equation 9 to 10, yields the update functions for $f(k, l)$, $g(k, l)$ and $z(k, l)$ in each iteration:

$$\begin{aligned} f_{k,l}^{n+1} &= f_{k,l}^n + \frac{\lambda}{1 - \lambda} [I_{k,l} - R(z_{k,l}^n, f_{k,l}^n, g_{k,l}^n)] \frac{\partial R}{\partial f} \Big|_{f_{k,l}} \\ g_{k,l}^{n+1} &= g_{k,l}^n + \frac{\lambda}{1 - \lambda} [I_{k,l} - R(z_{k,l}^n, f_{k,l}^n, g_{k,l}^n)] \frac{\partial R}{\partial g} \Big|_{g_{k,l}} \\ z_{k,l}^{n+1} &= z_{k,l}^n + \frac{\lambda}{1 - \lambda} [I_{k,l} - R(z_{k,l}^n, f_{k,l}^n, g_{k,l}^n)] \frac{\partial R}{\partial z} \Big|_{z_{k,l}} \end{aligned} \quad (11)$$

$f_{k,l}^n$, $g_{k,l}^n$ and $z_{k,l}^n$ are local 8-neighborhood mean of values around pixel position (k, l) . Thus the value for the $(n + 1)^{th}$ iteration can be estimated from the n^{th} iteration. During the iteration, the Lagrange multiplier λ will be gradually increased such that the smoothness constraint will be reduced as well. Since it is hard to estimate the initial values for z , we recompute z values after several iterations by integrating the normals using the method in [6].

Given initial values of f , g and z on the boundary, we can compute a numerical solution to the shape. Since variational problems usually depend on good initial guesses, we manually label the boundaries of each images. Results for simulation and real endoscopic images are showed in Fig.

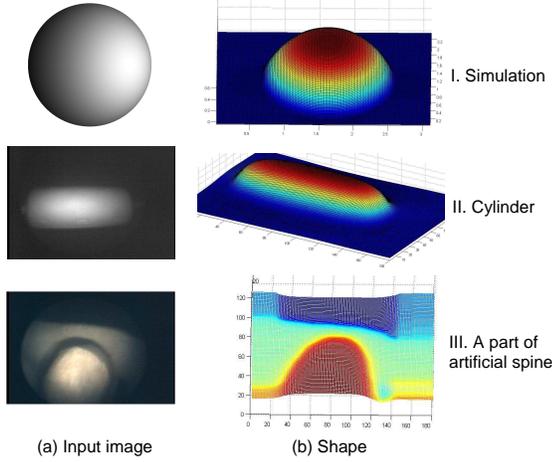


Figure 6. Results of shape from shading from single image. (a) input image. (b) shape from shading. I. synthesized sphere. II. cylinder. III. a part of spine.

6. This algorithm performs well for the synthesized image and endoscopic images of simple objects with complete boundaries. For the result of the artificial spine, however, a part of the shape is not accurate due to the partial boundary. Furthermore, due to the small field of view, the reconstruction from the single image cannot provide enough information about the complete shape of the bone.

4. Global Shape-from-Shading using Multiple Partial Views

As shown in Fig. 6, most images acquire a small fraction of visible contours. So, how do we estimate the complete shape reliably? An intuitive idea is to merge individual shapes recovered from different views. However, each image is defined in a view dependent local coordinates frame. We must first transform them to the same (world) coordinates before we can combine individual shapes. This is possible using a navigation tracking system. We attach optical markers on the endoscope (as showed in Fig. 3) and track it to obtain the transformation between the local and world coordinate frames for each image.

Consider a 3D-point set $S_{3 \times N_i}^i$ reconstructed from any image i , where N_i is the number of points in shape S^i . If all the shapes are reconstructed correctly, and without any error from the tracking system and the calibration process, they should be perfectly aligned in the world coordinates. Let ${}^wT_{m_i}$ denote the transformation from the marker at the position taking i^{th} image to the world coordinates, and A denote the intrinsic matrix of the camera that transforms the points in the marker coordinates to the camera coordinates (image distortion has been corrected in advance). Thus, any 3D-point set S^i in the local camera coordinates can be trans-

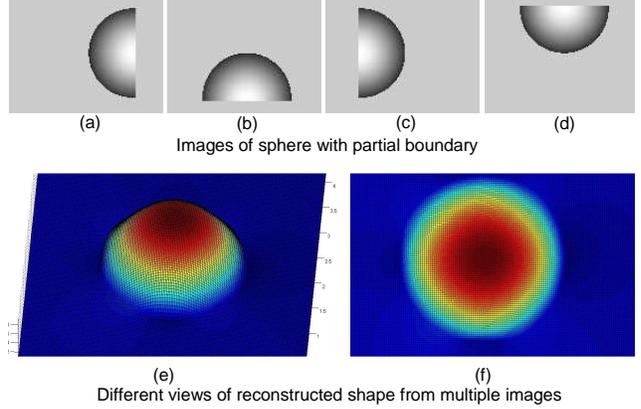


Figure 7. Simulation results of shape from shading from multiple views. (a)-(d) Synthesized images of different parts of a sphere. (e)-(f) Different views of reconstructed sphere.

formed to the world coordinates S_w^i by:

$$S_w^i = {}^wT_{m_i} \cdot [A^{-1} \cdot S^i; 1] \quad (12)$$

where, S_w^i is represented using homogenous coordinates in the world reference frame.

However, due to the tracking error (0.5mm) and calibration error (less than 6 pixels), S_w^i are not initially well matched (e.g. Fig. 8 (d), Fig. 9 (b)). Since the bone is rigid, we use ICP [1] to improve the matching between different views by ensuring overlapping fields of view in adjacent images. Then, for S_w^i , ICP yields an aligned shape ${}^{i-1}S_w^i$ with respect to the adjacent shape S_w^{i-1} :

$${}^{i-1}S_w^i = {}^{i-1}T_i * S_w^i \quad (13)$$

where, ${}^{i-1}T_i$ is the rigid transformation computed from ICP.

Using ICP, we can thus align all the shapes from one view with respect to the reference (say, first) view. In other words, the shape S_w^1 corresponds to the first image. Then, by combining Equation 12 and 13:

$${}^1S_w^i = \prod_{j=1}^{i-1} {}^jT_{j+1} \cdot {}^wT_{m_i} \cdot [M^{-1} \cdot S^i; 1] \quad (14)$$

where, ${}^1S_w^i$ is the aligned shape of S_w^i , with respect to the reference shape S_w^1 . Fig. 8 (e) (gray) and Fig. 9 (c) show the aligned shapes in the world coordinates, with respect to the first view. Contours or boundaries are aligned as well (See Fig. 8 (e) (color) and Fig. 9 (d)).

However, note that the partial shapes are only approximate. Thus, once the contours are aligned, we "re-grow" the shape in the local views but under the global constraints. For this, we compute the average value in Eq. 11 using the neighbors in the world coordinates. The key idea is to keep the global constraints updated according to the shape developed in the local view, and therefore SFS in the local views are guided towards convergence under global constraints.

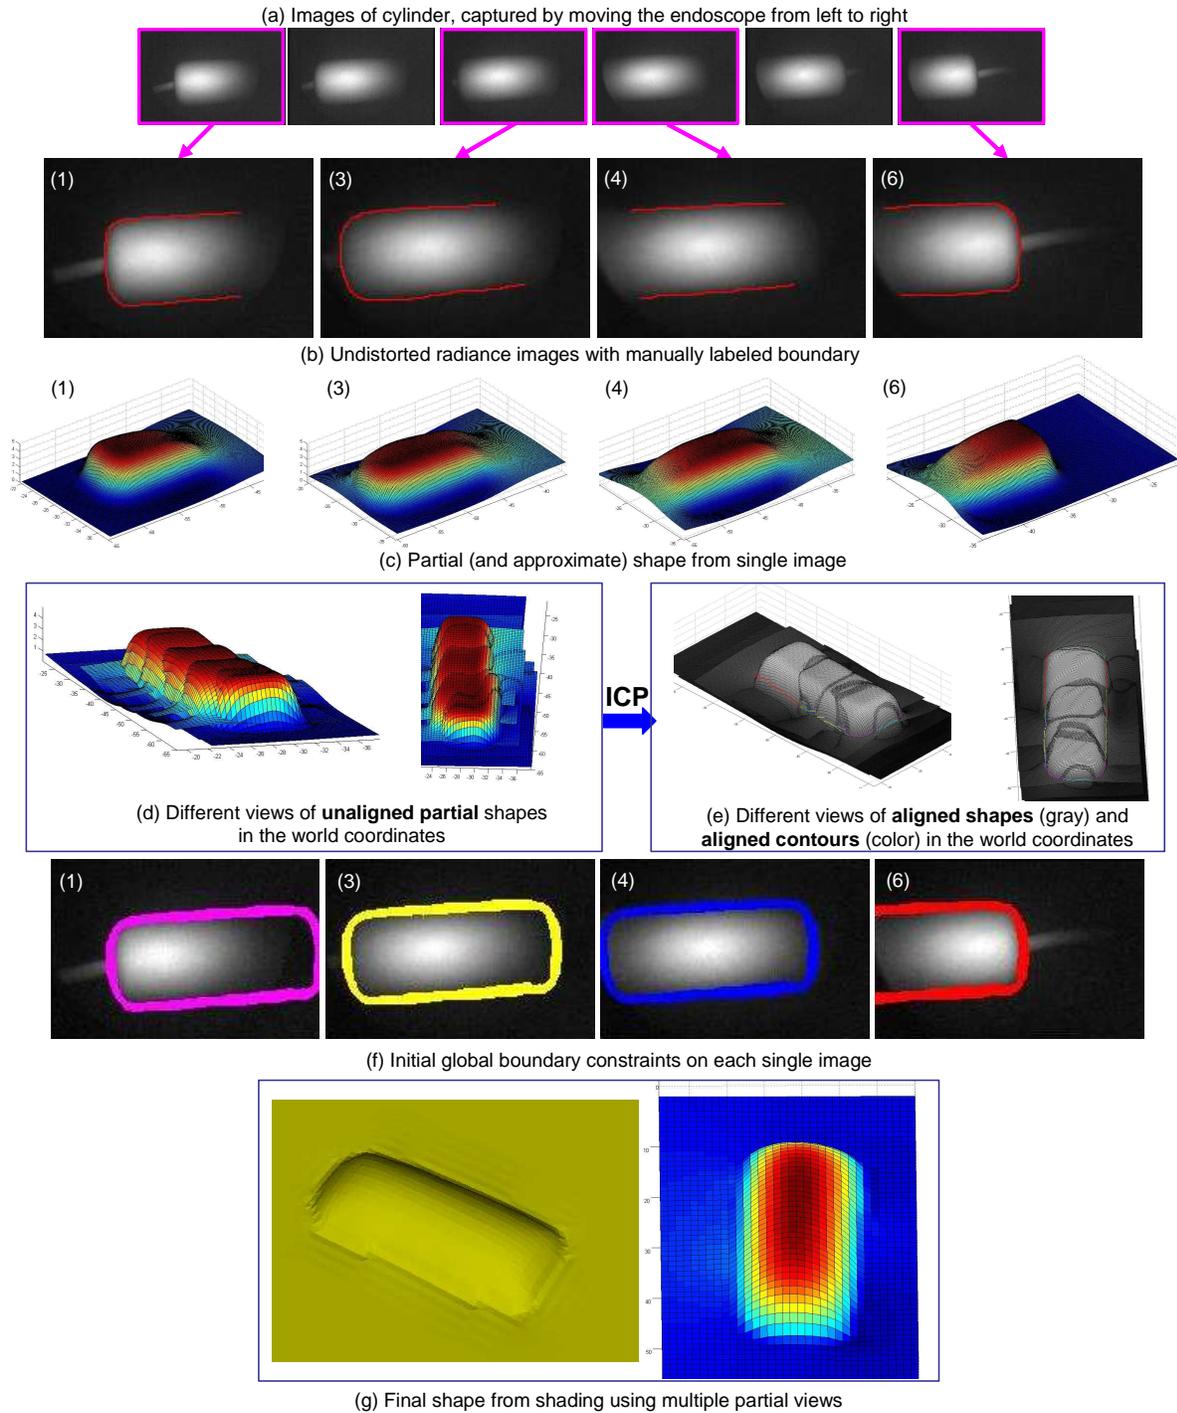


Figure 8. Reconstruction for the cylinder. (a) Six images are captured by moving the endoscope along the cylinder axis. Four of them are shown as an illustration. (b) After removing the distortion and illumination effects, the boundary in each image is labeled by hand, and then the initial normals are computed automatically. (c) Shape from each single image are reconstructed using the method described in Section 3. (d) Unaligned partial shapes in the world coordinates. (e) Aligned shapes (gray) and contours (color) using ICP in the world coordinates. (f) Updated local contours based on the global boundaries. (g) Final shape are reconstructed using the method described in Section 4. Note low resolution of global constraints can cause aliasing on the surface, which can be alleviated by increasing the resolution (see better result in Fig. 9 (e)).

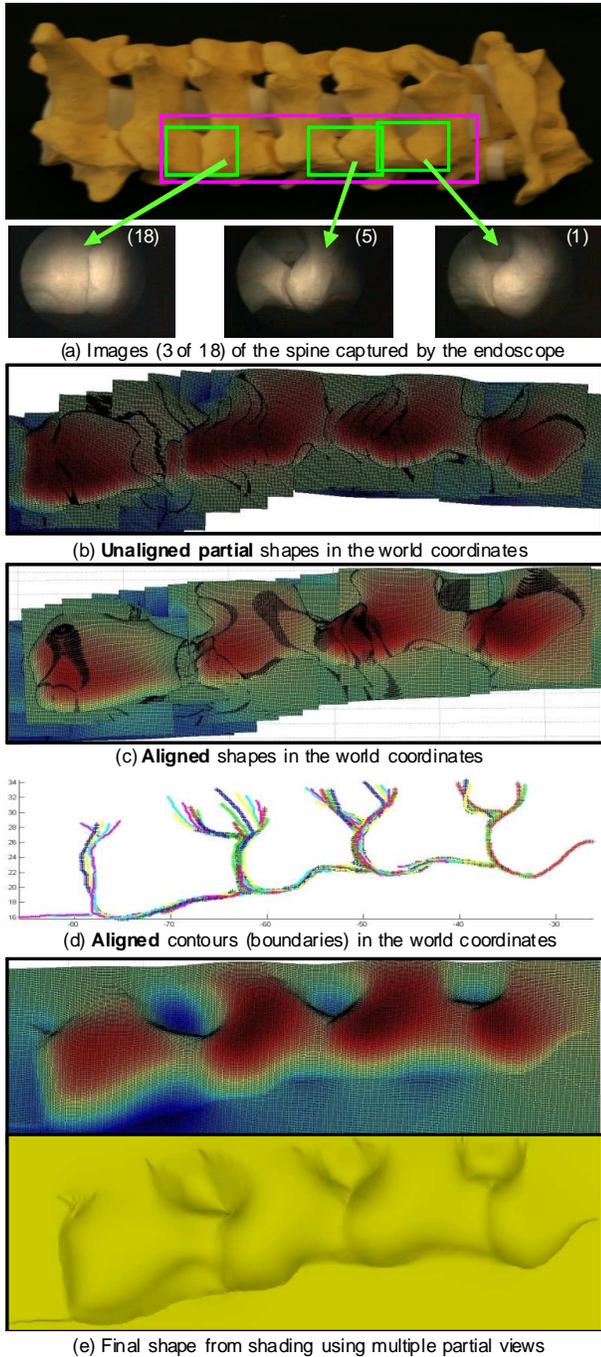


Figure 9. Reconstruction for the artificial spine. (a) 18 images are captured by moving the endoscope horizontally (only translation). (b) Unaligned partial partial shapes in the world coordinates. (c) Aligned shapes in the world coordinates. (d) Aligned contours (boundaries) in the world coordinates. (e) Final shape are reconstructed using the method described in Section 4.

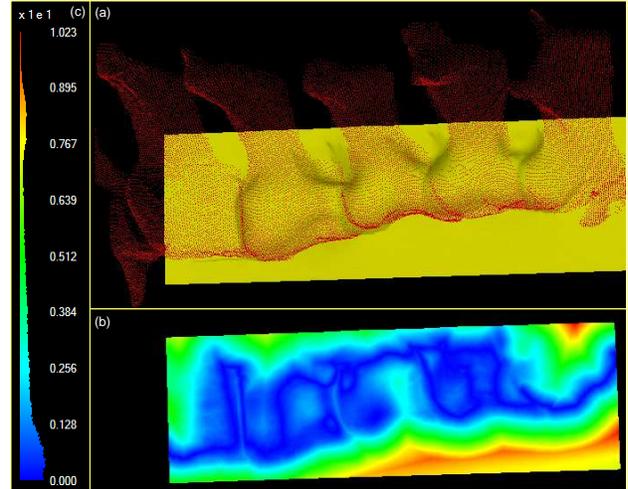


Figure 10. (a) Reconstructed shape (yellow) and ground truth surface (red). (b) Error map (vertex distance from the reconstructed shape to the ground truth surface) in HSV color space. (c) Error distribution.

We restart the SFS in the local coordinates frame simultaneously for all images. After initializing the normals and depth in each image, we can initialize the values on the corresponding global contours as well. If there are two or more points from different views that are mapped to the same point in the world coordinates, the normal and the depth at that point are computed as their average. In each iteration, we grow each image and update the global constraints. At the end of each iteration, we update the local normals according to the global normals. We proceed thus until the convergence. With global constraints, our algorithm converges quickly though not real-time (around 5 minutes for 18 images in Matlab, P4 2.4G CPU).

Fig. 7 shows the result of global shape-from-shading using four synthesized parts of a sphere. Fig. 8 shows the reconstructed shape of the cylinder (37mm x 12mm x 12mm) from 6 images. Fig. 9 shows the reconstruction of the spine (147mm x 60mm x 60mm) from 18 images. Please zoom-in to view the result in detail. Low resolution of the global constraints can cause aliasing as shown in Fig. 8 (g), which can be alleviated by increasing the resolution as shown in Fig. 9 (e).

Fig. 10 demonstrates the accuracy of our technique by comparing the reconstructed shape with the ground truth shape obtained using a laser range scanner. shown in Fig. 10. For comparison, we choose only the points that are on the surface of the spine. Then, the maximum, minimum, mean and RMS errors are 3.1mm, 0.0mm, 1.16mm and 1.5mm respectively. 10.23mm, 0.00mm, 2.82mm and 3.83mm respectively. This level of accuracy is practical for surgery.

5. Discussions and Conclusion

Shape-from-shading and shape-from-motion are both successfully used in many vision applications, but both have difficulty for orthopedic endoscopy due to the featureless bone surface and partial occluding boundaries in a small field of view. In this work, we propose a method to combine the strengths of both approaches to solve our problem: we formulate a shape-from-shading for single image under near point lighting and perspective projection, and develop a global shape-from-shading algorithm using multiple partial views. Due to the tracking and calibration error, we use ICP to improve the alignment of partial shapes and contours in the world coordinate frame. As a result, we obtain a reconstructed shape of a larger bone area providing useful visualization for surgical navigation in a minimally invasive procedure. There are several unsolved problems such as automatic boundary detection, obtaining good initial guesses, recovering high frequency details in shape, dealing with image noise due to blood and tissues. We hope to address these issues in the future.

Acknowledgements

This work is supported in parts by an NSF Grant #0325920 ITR: Data-driven Human Knee Modeling for Expert Surgical Planning Systems, an NSF CAREER Award #IIS-0643628, an NSF Award #CCF-0541307 and an ONR Award #N00014-05-1-0188.

References

- [1] P. Besl and N. McKay. A method for registration of 3-d shapes. *TPAMI*, 14(2):239–256, 1992. 5
- [2] M. J. Clarkson, D. Rueckert, A. P. King, P. J. Edwards, D. L. G. Hill, and D. J. Hawkes. Registration of video images to tomographic images by optimising mutual information using texture mapping. In *Proc. MICCAI'99*, volume LNCS 1679, pages 579–588, 1999. 1
- [3] K. Deguchi and T. Okatani. Shape reconstruction from an endoscope image by shape-from-shading technique for a point light source at the projection center. *Proc. MMBIA'96*, pages 290–298, 1996. 2, 4
- [4] D. Dey, D. G. Gobbi, P. J. Slomka, K. J. Surry, and T. M. Peters. Automatic fusion of freehand endoscopic brain images to three-dimensional surfaces: creating stereoscopic panoramas. *IEEE Trans. on Medical Imaging*, 21(1):23–30, 2002. 1
- [5] C. H. Q. Forster and C. L. Tozzi. Toward 3d reconstruction of endoscope images using shape from shading. *Proc. SIBGRAP'00, XIII Brazilian Symposium on Computer Graphics and Image Processing*, pages 90–96, 2000. 2, 4
- [6] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *TPAMI*, 10(4):439–451, 1988. 4
- [7] R. Hartley and A. Zisserman. *Multiple view geometry in Computer vision*. Cambridge University Press, 2nd edition, 2004. 2
- [8] B. K. P. Horn and M. J. Brooks. *Shape from Shading*. The MIT Press, Cambridge, Massachusetts, 1989. 1, 4
- [9] K. Ikeuchi and B. K. P. Horn. Numerical shape from shading and occluding boundaries. *AIJ*, 17(1):141–184, 1981. 4
- [10] A. K. S. Iyengar, H. Sugimoto, D. B. Smith, and M. S. Sacks. Dynamic in vitro quantification of bioprosthetic heart valve leaflet motion using structured light projection. *Annals of Biomedical Engineering*, 29:963–973, 2001. 2
- [11] R. Kimmel and J. A. Sethian. Optimal algorithm for shape from shading and path planning. *J. of Mathematical Imaging and Vision*, 14(3):237–244, 2001. 4
- [12] R. Kozera. An overview of the shape from shading problem. *Machine Graphics and Vision*, 7(1):291–312, 1998. 1
- [13] A. Litvinov and Y. Y. Schechner. Addressing radiometric nonidealities: a unified framework. *Proc. CVPR'05*, 2:52–59, 2005. 3
- [14] F. Mourgues, F. Devernay, G. Malandain, and E. Coste-Manière. 3d+t modeling of coronary artery tree from standard non simultaneous angiograms. *Proc. MICCAI'01*. 2
- [15] R. Onn and A. Bruckstein. Integrability disambiguates surface recovery in two-image photometric stereo. *Int. J. Comput. Vision*, 5(1):105, 1990. 4
- [16] C. Poelman and T. Kanade. A paraperspective factorization method for shape and motion recovery. *TPAMI*, 19(3):206–218, 1997. 2
- [17] M. Pollefeys, R. Koch, and L. VanGool. Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters. *IJCV*, 32(1):7–25, 1999. 2
- [18] E. Prados, F. Camilli, and O. Faugeras. A unifying and rigorous shape from shading method adapted to realistic data and applications. In *Journal of Mathematical Imaging and Vision*, volume 25, pages 307 – 328, 2006. 2
- [19] E. Prados and O. Faugeras. Perspective shape from shading and visibility solutions. In *Proc. ICCV'03*, volume 2, pages 826–831, 2003. 2, 4
- [20] E. Prados and O. Faugeras. Shape from shading: a well-posed problem? In *Proc. CVPR'05*, volume 2, pages 870–877, 2005. 2, 4
- [21] S. Seshamani, W. Lau, and G. Hager. Real-time endoscopic mosaicking. In *Proc. MICCAI'03*, volume LNCS 4190, pages 355–363, 2006. 2
- [22] D. Stoyanov, A. Darzi, and G. Z. Yang. A practical approach towards accurate dense 3d depth recovery for robotic laparoscopic surgery. *Computer Aided Surgery*, 10(4):199–208, 2005. 2
- [23] A. Tankus, N. Sochen, and Y. Yeshurun. Reconstruction of medical images by perspective shape-from-shading. *Proc. ICPR'04*, 3:778 – 781, 2004. 2
- [24] A. Tankus, N. Sochen, and Y. Yeshurun. Shape-from-shading under perspective projection. *IJCV*, 63(1):21–43, 2005. 2
- [25] R. Y. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IJRA*, pages 323–344, 1987. 3
- [26] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980. 2
- [27] C. Wu, S. G. Narasimhan, and B. Jaramaz. Endoscope calibration and derivation for shape from shading. Technical report, Carnegie Mellon University, 2007. 2, 3, 4
- [28] T. Yamaguchi, M. Nakamoto, Y. Sato, K. Konishi, M. Hashizume, N. Sugano, H. Yoshikawa, and S. Tamura. Development of a camera model and calibration procedure for oblique-viewing endoscopes. *Computer Aided Surgery*, 9(5):203–214, 2004. 3
- [29] R. Zhang, P. S. Tsai, J. E. Cryer, and M. Shah. Shape from shading: A survey. *TPAMI*, 21(8):690–706, 1999. 1
- [30] Z. Zhang. A flexible new technique for camera calibration. Technical Report MSR-TR-98-71, Microsoft Research, 1998. 3