ICCV03 Demonstation Proposal
3D Surface Acquisition from Single Images

Antonio Robles-Kelly[1*] and Edwin R. Hancock[2]
Department of Computer Science
The University of York
York, Y01 5DD, UK
[1]arobkell@cs.york.ac.uk [2]erh@cs.york.ac.uk

Abstract

Our aim in this demonstration proposal is to show how a surface acquisition system may be designed. We depart from the fact that a complete 3D representation of the object requires the depth information, the reflectance properties and the surface normals to be at hand. Following this rationale, our acquisition system makes use of shape-from-shading algorithms, surface height integration methods and BRDF approximation techniques for recovering the field of surface normals, the surface height and the reflectance properties of an object from a single image.

1 Demonstration Requirements

We require a table, two chairs, a power outlet, a projector with a screen and an A0-size board. In order to show the conference attendees how the 3D representation of a surface may be acquired from a single image, the presenters will provide a laptop computer and a poster.

2 Presenter Information

The presenter is one of the authors of this proposal and a member of the Computer Vision Group at the University of York. His contact details are

Antonio Robles-Kelly
Dept. of Computer Science
The University of York
Heslington, York YO105DD
United Kingdom
Tel. ++44(904)432774
Fax ++44(904)432767

*Corresponding author

3 Background

Surface acquisition provides a route to automatic object and scene modelling that is of potential practical use to the computer graphics community. Broadly speaking, existing methods may be divided into those that are geometrical in nature and those that are based upon photometric models. Turning our attention first to the geometric methods, the main contribution here has been to exploit projective geometry to develop algorithms for 3D object reconstruction [1]. Methodologically, the idea underpinning these algorithms is that of recovering both planar and curved surfaces from multiple views making use of calibration data and correspondance information. For curved surfaces there has been considerable success in using turntable sequences for surface reconstruction from both detected feature points and occluding contours [2]. Photometric methods, on the other hand, aim to recover surface information from shading or texture variations. This is a classical problem in the computer vision community, which has lead to a considerable literature on the development of algorithms for shape-from-shading [3], photometric stereo [4] and shape-from-texture [5]. However, the drawbacks of these algorithms are twofold. Firstly, they tend to oversmooth surface detail. Secondly, they are prone to error due to their numerical inestability.

Moreover, a complete 3D representation of the object requires not only the depth information, but also the surface normals and the BRDF to be at hand. This becomes evident ever since the normals and the BRDF are of capital importance to render the object, while the surface height data is essential when a mesh or spline representation is required. As a result, shape-from-shading algorithms alone are insufficient for acquiring the 3D representation of a surface from a single image. Following this rationale, our acquisition system makes use of shape-from-shading algorithms [3, 6] in conjunction with other computer vision methods for recovering the 3D representation of the object. We commence by recovering the field of surface normals (i.e. the Gauss map) making use of shape-from-shading methods. From the surface normals, surface height integration methods [7, 8] may be employed to recover the surface height. In parallel with this two step process, a variant of the BRDF approximation method of Robles-Kelly and Hancock [9] is used to recover the reflectance properties of the surface.

4 VisLab

Our demonstration vehicle is a program which we have named VisLab. VisLab is a piece of experimental software developed by the presenter that incorporates computer vision algorithms with and OpenGL-based visualisation GUI. VisLab allows the user to process image files (i.e. portable pixmaps, etc.) and propietary files containing OpenGL lists or fields of surface normals for the object. It also allows the use of what we have called “workspaces”, which comprise all the available data on the object (i.e. BRDF approximation, height data and Gauss map). In Figure 1, we show a sample view of VisLab’s main window with a document containing the image of a fragment of “The Three Graces” relief.

From a single image, the surface height data may be recovered making use of a wide variety of shape-from-shading and surface integration algorithms. It is well known that the performance of shape-from-shading algorithms may vary depending on the topology of the surface [10]. Hence, VisLab may process the surface making use of either of three shape-from-shading methods [11, 3, 6] and two surface integration algorithms [7, 8]. The user is free to choose the shape-from-shading algorithm and surface reconstruction method to be used and modify the parameters that may affect the quality of the recovered surface. We show this funtionality of our software in the video clips ”iccv_demo_01.mov” and ”iccv_demo_02.mov”.

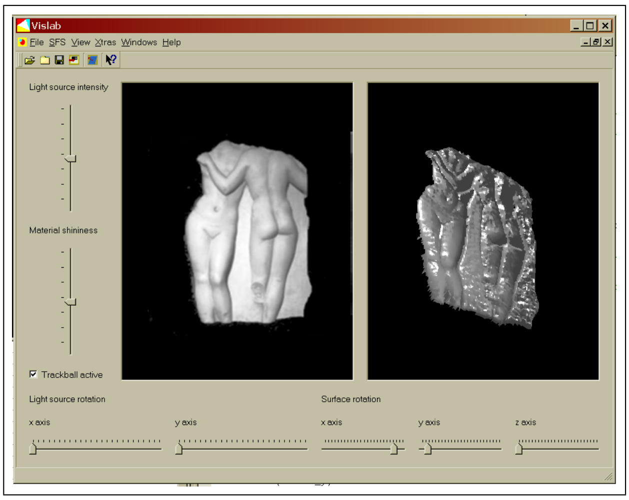

Figure 1: View of VisLab

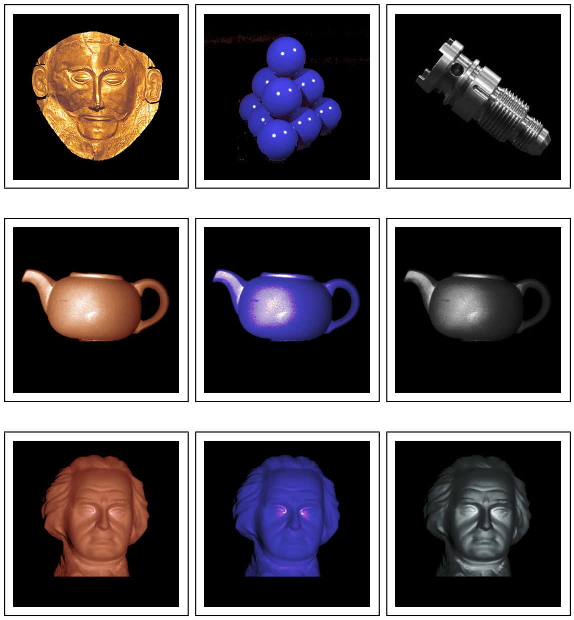

Figure 2: Top row: input images for the BRDF approximation process; Bottom rows: “teapot” and a bust of Betthoven rendered making use of the BRDF’s approximated from the input images.

It is worth noting that, since the field of normals, reflectance properties and height data are at hand, it is possible to perform image-based or forward rendering of the surface under study making use of the BRDF approximated from another open image or previously processed object. To illustrate this powerful feature of the software, we show some example results in Figure 2. In the top row, we show the real-world images used to estimate the BRDF’s. In the two bottom rows we show the results of performing forward rendering of a computer model of a teapot and a 3D-model of a bust of Beethoven acquired from a single gray-scale image using VisLab.

We will allow conference attendees to use VisLab under the supervision of the presenter. In this way, participants will be able to explore how visualisation tools like VisLab and recent vision research may be used to acquire realistic 3D representations from single gray-scale or color images.

5 Conclusions

Surface acquisition is of potential practical use to the computer graphics community. VisLab allows us to show how a practical system may be designed. We believe that the architecture of VisLab provides a means to combine efficiently the functionality of shape-from-shading methods, BRDF approximation techniques and surface integration to acquire the 3D representation.

References

[1] R. Fawcett, A. Zisserman, and J. M. Brady. Extracting structure from an affine view of a 3D point set with one or 2 bilateral symmetries. *Image and Vision Computing*, 12(9):615–622, November 1994.

[2] K. Wong, P. R. S. Mendonça, and R. Cipolla. Reconstruction and motion estimation from apparent contours under circular motion. In *Proc. of the British Machine Vision Conference*, pages I: 83–92, 1999.

[3] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(12):1250–1267, 1999.

[4] F. Bernardini and H. Rushmeier. High-quality texture reconstruction from multiple scans. *IEEE Trans. on Visualization and Comp. Graphics*, 7(4):318–332, 2001.

[5] B. J. Super and A. C. Bovik. Planar surface orientation from texture spatial-frequencies. *Pattern Recognition*, 28(5):729–743, May 1995.

[6] A. Robles-Kelly and E. R. Hancock. A mumford-shah diffusion process for shape-from-shading. In *Proceedings of the British Machine Vision Conference*, pages 708–717, 2002.

[7] A. Robles-Kelly and E. R. Hancock. A graph-spectral method for surface height recovery from needle maps. In *Proc. of the IEEE Conference on Conputer Vision and Pattern Recognition*, pages I:141–148, 2001.

[8] A. G. Bors, E. R. Hancock, and R. C. Wilson. Terrain analysis using radar imagery. To appear in *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2003.

[9] A. Robles-Kelly and E. R. Hancock. Lambertian correction for rough and specular surfaces. In *Proceeding of Video, Vision and Graphics*, pages 213–220, 2003.

[10] R. Zhang, Ping-Sing Tsai, J. E. Cryer, and M. Shah. Shape from shading: A survery. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999.

[11] M. Bichsel and A. P. Pentland. A simple algorithm for shape from shading. In *CVPR92*, pages 459–465, 1992.