Projecting Surface Curvature Maps

John Rugis,[*] Department of Computer Science, University of Auckland

Figure 1: Michelangelo's David: (i) triangle mesh (flat), (ii) triangle mesh (smooth), (iii) with a strip of projected curvature.

1 Introduction

High resolution 3D range scan digitizations often produce data sets that, when directly converted to 3D models, are too large to be rendered quickly for display. The usual solution to this problem is to sacrifice detail and reduce the model size by simplification. In this research, we propose a method whereby high resolution detail can be merged with a simplified mesh model for visualization. To this end, we firstly introduce curvature maps which are calculated from the original scan data associated with each scan view point. Then we project these curvature maps onto a simplified model in a way that is analogous to using slide projectors to project multiple photographs onto a 3D object.

There are a number of problems associated with this approach that needed to be solved. Generally raw scan data contains noise, and this noise destroys the possibility of direct curvature calculation. Also, the standard 3D graphics lighting projection model is not appropriate because it does not handle the unavoidable projection overlap seamlessly.

2 Curvature Maps and Projection

When working with point set data, surface curvature can only be estimated [Klette and Rosenfeld 2004]. In this research, we use an uncompensated orthogonal cut method to calculate a mean curvature as described in [Rugis 2005]. We handle the noisy data problem by firstly doing a shading encoded mean curvature calculation on the noisy data and mapping the noisy curvatures into the 2D domain. Standard 2D image filtering and segmentation is applied before mapping the results back into vertex colors in the 3D domain. Each perspective correct final projection map consists of a single frame vertex colored rendering of these results.

A goal from the outset was to, as much as possible, use existing 3D graphics techniques. However, the standard lighting projection model is additive for multiple sources [Foley et al. 1996]. We introduce a multiple source averaging lighting model that overcomes this problem. The diffuse component for each of the projection sources is averaged into a total value that is used to modulate the final standard lighting value.

3 Experiments and Conclusion

We tested our approach on the David dataset from the Digital Michelangelo Project [Levoy et al. 2000]. Resultant renderings of single image frames are shown in Figure 1. Rendered images that specifically illustrate the success of our averaging lighting model are in an accompanying file (lighting_models.jpg). Because the curvature maps were created outside any lighting model and the fact that curvature itself is rotation and translation invariant, the projected surface curvature detail retains a unique clarity in animated visualization. (See the accompanying video.)

References

FOLEY, J., VANDAM, A., FEINER, S., AND HUGHES, J. 1996. *Computer Graphics: Principles and Practice*. Addison-Wesley, Boston.

KLETTE, R., AND ROSENFELD, A. 2004. *Digital Geometry*. Morgan Kaufmann, San Francisco.

LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINZTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. 2000. The digital michelangelo project. In *Proceedings of SIGGRAPH 2000*, ACM Press / ACM SIGGRAPH, K. Akeley, Ed., 131–144.

RUGIS, J. 2005. Surface curvature maps and Michelangelo's David. In *Image and Vision Computing New Zealand 2005*, B. McCane, Ed., 218–222.

[*]e-mail: john.rugis@manukau.ac.nz