CS590G: Final Project – Create your own capture/model/render system

Out: October 28, 2004
Due: December 9, 2004

Objective

The goal of the final project assignment is to allow you to use your imagination to design, implement and experience your own capture system. The previous assignments have helped you to incrementally build up a complete system. For your project, you may either build upon the previous assignments and extend the system in a particular direction or implement a new system.

Each project will be performed by small teams of 1 to 3 people, all taking CS590G this semester. As part of the assignment, you will each need to present a relevant research paper, the group will present a project proposal, a mid-project demo, and during demo day at end of semester each group will give an exciting demo and presentation of their work. During the project proposal and demos, you must identify which component of the system each team member did. As with previous assignments, I much prefer a system that works well and fewer bells and whistles than a more complex system that does not work well.

Below is a list of projects. You may use these as starting points for your project proposal. Nevertheless, you may also present a completely new project proposal. All project proposals are pending my approval.

Capture

Camera Design Program. Create a design program to allow you to visualize a (synthetic) environment as seen by a simulated camera. This program should support camera and mirror arrangements for standard FOV and large FOV cameras.

- "A Theory of Single-Viewpoint Catadioptric Image Formation", Simon Baker and Shree K. Nayar, International Journal of Computer Vision, 1999.
- "A True Omnidirectional Viewer", Vic Nalwa, Technical report, Bell Laboratories, Holmdel, NJ 07733, USA.
- Panoramic Vision, Ryad Benosman, Sing Bing Kang, Olivier Faugeras, Springer Verlag, 1st edition, 2001.

Robust and Portable Camera Pose Estimation System. Create a system that allows you to quickly setup a scenario for fast but relatively accurate camera pose estimation. For example, deploy landmarks and from their projections compute pose. I have some initial code you can use as a starting point that 'works' but the objective is too significantly improve its robustness and to enable a wider range of pose estimation. Pose estimation can be limited to certain areas, e.g. a table top, a room, etc. Let your imagination go wild here.

- Superior Augmented Reality Registration by Integrating Landmark Tracking and Magnetic Tracking, Andrei State, Gentaro Hirota, David T. Chen, William F. Garrett and Mark A. Livingston. In *Computer Graphics* Proceedings, Annual Conference Series, ACM SIGGRAPH, pp. 429-438, 1996.

Capture Device Survey. Write a survey or give a complete presentation on a taxonomy of multiple capture devices and how they capture, model, and render environments (maybe not relevant enough to this course ???).

Static Modeling and Rendering

Automatic City Modeling. Generate a city model using a set of urban and architectural rules (ideally rules obtained from a real city or neighborhood).

- Procedural Modeling of Cities, Yoav Parish, Pascal Muller, Proceedings of ACM SIGGRAPH, pp. 301-308, 2001.
- Instant Architecture, Peter Wonka (Georgia Institute of Technology), Michael Wimmer (Vienna University of Technology), François Sillion (ARTIS/INRIA Rhône-Alpes), Bill Ribarsky (Georgia Institute of Technology), Proceedings of ACM SIGGRAPH, 2003.

Self-calibration and Reconstruction. Exploiting properties of epipolar geometry and laws of physics simultaneously do camera calibration and object reconstruction. You may also view this as “bypassing” the calibration phase and going straight to the reconstruction phase.

Lots of papers…

4D Plenoptic Function Renderer with Dynamic Reparameterization. Capture many images so as to sample a 4D plenoptic function. Use (a) a classical two-plane parameterization of (e.g., Lightfield/Lumigraph), (b) a viewpoint-centric parameterization (e.g., Sea-of-Images, Plenoptic Stitching), (c) make your-own parameterization. To help interpolate images use (a) more images (e.g., Lightfield), (b) a proxy (e.g., Lumigraph), (c) feature correspondences -- manually determined is ok (e.g., Sea of Images). Render novel views by extracting individual pixels from images or via image blending and image morphing. In addition, allow the images to be reparameterized so as to support apparent aperture changes and focal length changes.

- Marc Levoy and Patrick M. Hanrahan, Proceedings of SIGGRAPH 96, pages 31–42, August 1996.
- The Lumigraph, Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen, Proceedings of ACM SIGGRAPH 1996, pages 43–54, August 1996.
- "Interactive Image-Based Rendering Using Feature Globalization", Daniel G. Aliaga, Dimah Yanovsky, Thomas Funkhouser, Ingrid Carlbom, Proceedings of ACM Symposium on Interactive 3D Graphics, April, 2003.

- Dynamically Reparameterized Lightfields, Aaron Isaksen, Leonard McMillan, Steven Gortler, Proceedings of ACM SIGGRAPH, 2000.

Extract an Optimal 3D Point Set from the Features. Given a set of images and corresponded features (which may be determined manually), estimate a 3D position for each observed feature. Back project the feature onto each image and compute an error value. Perturb the 3D estimation until achieving a single optimum location for all images. Then, optionally, iterate between perturbing the 3D coordinates and the camera poses. This project may also be interpreted as computing the "depth" to a feature from n-images (as opposed to just 2 images as would be the case with classical stereo depth estimation).

- Bundle adjustment - a modern synthesis. B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. In *Vision Algoritms: Theory andPractic*e. Springer-Verlag, 2000.
- A Maximum-Flow Formulation of the N-camera Stereo Correspondence Problem**,** Sébastien Roy and Ingemar J. Cox, *Int. Conf. on Computer Vision (ICCV'98)*, Bombay, India, January 1998, p. 492-499.
- Three Dimensional Computer Vision, Olivier Faugeras, MIT Press, 1996.

Build a 3D Surface Model. In addition to the above, fit a surface though all the points.

- Creating Generative Models from Range Images, Ravi Ramamoorthi, James Arvo, Proceedings of ACM SIGGRAPH, 1999.

Build a 3D Volumetric model. From the array of images, compute the contents of a volume of space in front of the cameras. Use a space-carving, voxel-coloring, or a new approach of your own.

- Photorealistic Scene Reconstruction by Voxel Coloring. S. M. Seitz and C. R. Dyer, International Journal of Computer Vision, 35(2), 1999, pp. 151-173. Shorter version in Proc. Computer Vision and Pattern Recognition Conf., 1997, 1067-1073.
- A Theory of Shape by Space Carving, K. N. Kutulakos and S. M. Seitz, International Journal of Computer Vision, Marr Prize Special Issue, 2000, 38(3), pp. 199-218. Earlier version appeared in Proc. Seventh International Conference on Computer Vision (ICCV), 1999, pp. 307-314.

Create a Model using a Level-Set Approach. Implement a level-set based method to extract the surface.

- Geometric Level Set Methods in Imaging, Vision and Graphics, S. Osher and N. Paragios (eds), Springer Verlag, August 2002.

Compute a BRDF from the Image Dataset. Once you have a 3D model, compute multiple values per pixel, indexed by incoming/outgoing ray direction. This builds up a BRDF for the patch. This would then allow re-illumination, for instance.

Lots of papers…

Dynamic Modeling and Rendering

5D Plenoptic Function. Similar to the 4D Plenoptic Function Renderer but capture a rigid object undergoing simple periodic motion, e.g. rotating about an axis. Allow for novel views but also allow for selection of time parameter.

Live Stereo Viewing. View a scene that you have already (mostly) reconstructed but generate stereo views (e.g. for viewing using glasses) in real-time. This requires slightly changing the view for one eye. Really cool!

- Hybrid Stereo Camera: An IBR Approach for Synthesis of Very-High-Resolution Stereoscopic Image Sequences, Harpreet S. Sawhney, Yanlin Guo, Keith Hanna, Rakesh Kumar, Sean Adkins, Samuel Zhou, Proceedings of ACM SIGRAPH 2001.

Additional Papers

Image-Based Registration and Reconstruction Systems

Realistic surface reconstruction of 3D scenes from uncalibrated image sequences, Koch, M. Pollefeys, L. Van Gool, Journal Visualization and Computer Animation, Vol. 11, pp. 115-127, 2000.

Real-Time 3D Model Acquisition, Szymon Rusinkiewicz, Olaf Hall-Holt, Marc Levoy, Transactions on Graphics (SIGGRAPH proceedings), 2002.

Calibrated, Registered Images of an Extended Urban Area, Seth Teller, Matthew Antone, Zachary Bodnar, Michael Bosse, Satyan Coorg, Manish Jethwa, and Neel Master, International Journal of Computer Vision, 2003.

Spacetime Stereo: Shape Recovery for Dynamic Scenes, L. Zhang, B. Curless, and S. M. Seitz, Proc. Computer Vision and Pattern Recognition Conf. (CVPR), 2003.

From 2D Images to 3D Space

Photorealistic Scene Reconstruction by Voxel Coloring. S. M. Seitz and C. R. Dyer, International Journal of Computer Vision, 35(2), 1999, pp. 151-173. Shorter version in Proc. Computer Vision and Pattern Recognition Conf., 1997, 1067-1073.

A Theory of Shape by Space Carving, K. N. Kutulakos and S. M. Seitz, International Journal of Computer Vision, Marr Prize Special Issue, 2000, 38(3), pp. 199-218. Earlier version appeared in Proc. Seventh International Conference on Computer Vision (ICCV), 1999, pp. 307-314.

Image-Based Visual Hulls, Matusik, Wojciech, Chris Buehler, Ramesh Raskar, Steven J. Gortler, and Leonard McMillan, Proceedings of ACM SIGGRAPH 2000, (New Orleans, LA July 23-28, 2000), pp. 369-374.

Image-Based 3D Photography Using Opacity Hulls, Matusik, Wojciech, Hanspeter Pfister, Addy Ngan, Paul Beardsley, and Leonard McMillan, Proceedings of ACM SIGGRAPH 2002, (in Computer Graphics Proceedings, Annual Conference Series, July 2002), pp. 427-437.

Generative Models

Creating Generative Models from Range Images, Ravi Ramamoorthi, James Arvo, Proceedings of ACM SIGGRAPH, 1999.

Procedural Modeling of Cities, Yoav Parish, Pascal Muller, Proceedings of ACM SIGGRAPH, pp. 301-308, 2001.

Instant Architecture, Peter Wonka (Georgia Institute of Technology), Michael Wimmer (Vienna University of Technology), François Sillion (ARTIS/INRIA Rhône-Alpes), Bill Ribarsky (Georgia Institute of Technology), Proceedings of ACM SIGGRAPH, 2003.

Systems for Managing Large Models

Thomas A. Funkhouser and Carlo H. Sequin. Adaptive Display Algorithms for Interactive Frame Rates During Visualization of Complex Virtual Environments. Computer Graphics (SIGGRAPH `93), Los Angeles, CA, August, 1993, p. 247-254.

Hierarchical Image Caching for Accelerated Walkthroughs of Complex Environments, Jonathan Shade, Dani Lischinski, David H. Salesin, Tony DeRose, and John Snyder, Proceedings of ACM SIGGRAPH, 1996.

Miscellaneous

Omnivergent Stereo. H.Y. Shum, A. Kalai, and S. M. Seitz, Proc. Seventh International Conference on Computer Vision (ICCV), 1999.

Plenoptic Sampling, Jin-Xiang Chai, Shing-Chow Chan, Heung-Yeung Shum, Xin Tong. Proceedings of ACM SIGGRAPH, pp. 307-318, 2000.

Superior Augmented Reality Registration by Integrating Landmark Tracking and Magnetic Tracking, Andrei State, Gentaro Hirota, David T. Chen, William F. Garrett and Mark A. Livingston. In *Computer Graphics* Proceedings, Annual Conference Series, ACM SIGGRAPH, pp. 429-438, 1996.

Conferences and Journals

You may also select papers/articles from the following conferences: SIGGRAPH (www.cs.brown.edu/~tor) , I3D: Symposium on Interactive 3D Graphics, 3DIM: 3D Digital Imaging and Modeling, Vis: IEEE Visualization, GI: Graphics Interface, Eurographics (and workshops), CGI: Computer Graphics International, VRAIS: Virtual Reality Annual International Symposium, and VR: IEEE Virtual Reality.

And the following journals: IEEE Transactions on Visualization and Computer Graphics, Graphics, IEEE Transactions on Image Processing, IEEE Transactions on Circuits and System for Video Technology (has some very relevant articles).

There are other acceptable conferences and journals. *Email me to obtain approval if you wish to use material from another source.*