Organizer
Radek Grzeszczuk
Intel Corporation

Lecturers
Jean-Yves Bouguet
Radek Grzeszczuk
Intel Corporation

Leonard McMillan
Massachusetts Institute of Technology

Ko Nishino
University of Tokyo

Hanspeter Pfister
Mitsubishi Electric Research Lab

Marc Pollefeys
Katholieke Universiteit Leuven

44

Image-Based Modeling

Course Notes

The 29th International Conference on
Computer Graphics and Interactive Techniques

Henry B. Gonzalez Convention Center
San Antonio, Texas USA

Conference: 21–26 July 2002
Exhibition: 23–25 July 2002

SAN ANTONIO
SIGGRAPH
2002

Image-Based Modeling

SIGGRAPH 2002 Course 44 Notes

Organizer

Radek Grzeszczuk
Intel Corporation

Lecturers

Marc Pollefeys
Katholieke Universiteit Leuven

Leonard McMillan
Massachusetts Institute of Technology

Hanspeter Pfister
Mitsubishi Electric Research Lab

Ko Nishino
University of Tokyo

Jean-Yves Bouguet
Intel Corporation

Radek Grzeszczuk
Intel Corporation

Abstract

As the complexity of image-based models grows, researchers are facing an increasingly challenging problem of collecting and processing massive amounts of radiance data. Nowadays image-based data sets often consist of tens of thousands of images or hundreds of gigabytes. Clearly, acquisition of these data requires new methodologies. New techniques must also be developed for compact representation and efficient visualization of resulting image-based models. This course is an overview of new technologies for collecting and analyzing densely sampled radiance data and building image-based models that are compact, accurate and easy to render.

Topics include geometry reconstruction, registration of image data with geometry, acquisition of radiance data, data-driven compression of radiance data, hardware-accelerated rendering of image-based models, image-based models for computer games and visualization.

Contents

Session 5: Image-Based 3D Photography using Opacity Hulls (Part I) — Leonard McMillan

Slides: "Image-Based 3D Photography using Opacity Hulls"
L. McMillan, H. Pfister
Paper: "Image-Based Visual Hulls"
W. Matusik, C. Buehler, R. Raskar, S. Gortler, L. McMillan

Session 6: Image-Based 3D Photography using Opacity Hulls (Part II) — Hanspeter Pfister

Paper: "Surface Splatting"
M. Zwicker, H. Pfister, J. van Baar, M. Gross
Paper: "Image-Based 3D Photography using Opacity Hulls"
W. Matusik, H. Pfister, A. Ngan, P. Beardsley, R. Ziegler, L. McMillan

Lecturer Biographies

Marc Pollefeys is a post-doctoral researcher at the Center for Processing of Speech and Images of the K.U.Leuven. His PhD dissertation on "Self-calibration and Metric 3D Reconstruction from Uncalibrated Image Sequences" was awarded the Scientific Prize BARCO 1999. His current research focuses on 3D modeling from images, multi-view geometry, image-based rendering, virtual and augmented reality and applications. He is currently involved in projects ranging from digital archaeology, over 3D TV, to planetary rover control. Marc Pollefeys has written over 60 technical papers and won several awards, amongst which the prestigious Marr Prize at the International Conference on Computer Vision in 1998. He has organized the course on "Obtaining 3D models with a hand-held camera" at SIGGRAPH 2000 and 2001 and similar courses at major vision conferences and has also contributed to the previous edition of this course.

Leonard McMillan is a pioneer in the area of image-based rendering. Image-based rendering (IBR) is a new approach to computer graphics in which scenes are modeled using a collection of reference images. These reference images can then be used to synthesize new renderings from a wide range of viewing positions. He has worked a wide range of different approaches to IBR including warping images with depth, light field rendering, and generating view-dependent models directly from live video streams. Leonard is also interested in a wide range of related topics including three-dimension display technologies, computer graphics hardware, and the fusion of image processing, multimedia, and computer graphics.

Leonard is an Associate Professor in the EECS Department and a member of the Computer Graphics Group of the Laboratory for Computer Science at MIT. Leonard received his BSEE and MSEE from Georgia Institute of Technology and his Ph.D. from the University of North Carolina at Chapel Hill. He has also worked at Bell Laboratories and Sun Microsystems.

Hanspeter Pfister is Associate Director and Senior Research Scientist at MERL - Mitsubishi Electric Research Laboratories - in Cambridge, MA. He is the chief architect of VolumePro, Mitsubishi Electric's real-time volume rendering hardware for PCs. His research interests include computer graphics, scientific visualization, and computer architecture. His work spans a range of topics, including point-based rendering and modeling, 3D scanning and 3D photography, and computer graphics hardware. Hanspeter Pfister received his Ph.D. in Computer Science in 1996 from the State University of New York at Stony Brook. He received his M.S. in Electrical Engineering from the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, in 1991. Dr. Pfister has taught courses at major graphics conferences including SIGGRAPH, IEEE Visualization, and Eurographics. He is Associate Editor of the IEEE Transactions on Visualization and Computer Graphics (TVCG), member of the Executive Committee of the IEEE Technical Committee on Graphics and Visualization (TCVG), and has served as a member of international program committees of major graphics conferences. Dr. Pfister is the general chair of the IEEE Visualization 2002 conference in Boston. He is member of the ACM, ACM SIGGRAPH, IEEE, the IEEE Computer Society, and the Eurographics Association.

Ko Nishino is a postdoctoral researcher at the Institute of Industrial Science, University of Tokyo and Japan Science and Technology Corp. Ko Nishino received his Ph.D. in Information Science in March 2002 from The University of Tokyo. He received his M.S. degree in Electrical Engineering from The University of Tokyo in 1999. He has been working on research topics in computer vision and computer graphics with special interest in image-based rendering and modeling, inverse-rendering, physics-based vision and color. His recent work in these areas, including his dissertation, has focused on efficient representations of real world objects given a dense/sparse set of images and a geometric model. He has also worked on geometric modeling, including registration and integration of range images and 2D-3D alignment. His work has been published at International Conference on Computer Vision (ICCV) and Computer Vision and Pattern Recognition (CVPR). He is a member of ACM, IEEE, the IEEE Computer Society.

Jean-Yves Bouguet is a Senior Researcher at Intel's Microprocessor Research Labs since 1999. He received his diplome d'ingenieur from the Ecole Superieure d'Ingenieurs en Electrotechnique et Electronique (ESIEE, Paris) in 1994 and the M.S. and Ph.D. degrees in Electrical Engineering from the California Institute of Technology (Caltech) in 1994 and 1999, respectively. His research interests cover all computer vision techniques (passive and active) for capturing the three-dimensional structure of real scenes. During his thesis work, he has developed a simple and inexpensive method for scanning objects using shadows. This work was first presented at ICCV'98 and a patent is pending on that invention. Recently, Jean-Yves' work focused on developing modeling techniques that combine 3D geometry capture and scene reflectance acquisition for realistic rendering of real and synthetic scenes with complex shape and surface characteristics. Jean-Yves has received a number of distinctive awards including the J. Walker von Brimer award for "extraordinary accomplishments in the field of 3D photography". He also collaborated with Prof. Jim Arvo, Prof. Peter Schroder and Prof. Pietro Perona in teaching a graduate level course on 3D photography from 1996 to 1998 at Caltech. Jean-Yves also contributed to the 3D photography course offered at SIGGRAPH 1999 and 2000 with Prof. Brian Curless, Prof. Steve Seitz, Dr Paul Debevec, Prof. Marc Levoy and Prof. Shree Nayar.

Radek Grzeszczuk joined Intel's Microprocessor Research Labs in 1998 as a Senior Researcher. He received his Ph.D. degree (1998) and his M.S. degree (1994) in Computer Science from the University of Toronto. His PhD thesis research was done under the supervision of Demetri Terzopoulos and Geoffrey Hinton and focused on using neural networks for fast emulation and control of physics-based models. The results of this work were published at SIGGRAPH'98 and NIPS'98. Radek Grzeszczuk's pioneering research with Stephen Gortler, Michael Cohen and Richard Szeliski at Microsoft Research Graphics Group on image-based rendering culminated in the publication of "The Lumigraph" at SIGGRAPH'96. His recent work on image-based modeling and rendering focuses on methods for efficient representation and visualization of complex shape and reflectance properties of objects. He published a number of important scientific papers, primarily in computer graphics, but also in artificial life, neural networks, and computer vision. He received an award in 1995 from Ars Electronica, the premier competition for creative work with digital media, for his work on artificial animals for computer animation and virtual reality.

Lecturer Contact Information:

Marc Pollefeys, Postdoctoral Researcher, Katholieke Universiteit Leuven
Marc.Pollefeys@esat.kuleuven.ac.be
http://www.esat.kuleuven.ac.be/~pollefey

Leonard McMillan, Assistant Professor, Massachusetts Institute of Technology
mcmillan@graphics.lcs.mit.edu
http://www.graphics.lcs.mit.edu/~mcmillan/

Hanspeter Pfister, Associate Director and Research Scientist, Mitsubishi Electric Research Lab
pfister@merl.com http://www.merl.com/people/pfister/index.html

Ko Nishino, Postdoctoral Researcher, University of Tokyo
kon@cvl.iis.u-tokyo.ac.jp
http://www.cvl.iis.u-tokyo.ac.jp/~kon

Jean-Yves Bouguet, Senior Researcher, Intel Corporation
Jean-Yves.Bouguet@intel.com
http://www.intel.com/research/mrl/people/bouguet_j.htm

Radek Grzeszczuk, Senior Researcher, Intel Corporation
Radek.Grzeszczuk@intel.com
http://www.intel.com/research/mrl/people/grzeszczuk_r.htm

Course Introduction and Overview
Radek Grzeszczuk

The techniques discussed in this course suggest that the easiest and the most efficient method of representing radiance properties of an object is in the form of a sampled function parameterized directly on the geometry of the object. The underlying assumption is that this function is computed from discrete examples of the radiance data available in the form of images. There are many facts that substantiate the approach proposed in this course. Historically, graphics hardware has been optimized for rendering of geometrical models, therefore geometry-based approach guarantees efficiency. Additionally, the representation of radiance data directly on the geometry of the object leads to higher compression rates and less artifacts than purely image-based representations.

There is also strong justification for sample-based representation of radiance data. Realistic analytical radiance models are difficult to develop and render efficiently. However, radiance data can be easily sampled from physical world through image acquisition. Although it is possible to fit predefined analytic models into the sampled radiance data, the result is often not photorealistic and potentially inefficient to render. A simple way of representing complex radiance data is in the form of a sampled function, such as a texture map. This representation is not only more natural, it should also be more efficient, since texture mapping support is offered on all the existing graphics hardware.

As the complexity of image-based models grows, researchers are facing an increasingly challenging problem of collecting and processing massive amounts of radiance data. Nowadays image-based data sets often consist of tens of thousands of images or hundreds of gigabytes. Clearly, acquisition of these data requires new methodologies. New techniques must also be developed for compact representation and efficient visualization of resulting image-based models. Consequently, data-driven dimensionality reduction methods for analysis of radiance data are becoming increasingly important. This course is an overview of new technologies for collecting and analyzing densely sampled radiance data and building image-based models that are compact, accurate and easy to render.

The topics covered in this course can be divided into two main themes: acquisition of image-based models and modeling and rendering of image-based models. The course will progress according to the six sessions summarized below:

1. **Re-rendering from a Dense/Sparse Set of Images (Ko Nishino):** This segment will show two different methods to efficiently represent the radiance data of real objects. First, I will overview an appearance-based approach to efficiently represent the radiance data. This method densely samples appearance variation of a real object under various illumination and viewing conditions, and compresses it in the 2D coordinate system defined on the 3D model surface. By handling radiance data on the object surface, the method successfully boosts the correlation of its variation and consequently achieves more efficient representa-

tion compared to pure image-based methods. Next, I will show a model-based approach for extraction and representation of photometric properties of real objects. Unlike previous model-based approaches, we recover all three unknowns, namely diffuse texture, parametric BRDF model and lighting, from a sparsely sampled radiance data. By approximating the illumination distribution as a point light source set on the surface of a hemisphere, we formalize specular reflection as convolution on that surface and accomplish blind deconvolution for simultaneous estimation of lighting and reflectance parameters. Both methods will be explained with results using real objects.

2. **Acquisition of Surface Light Fields (Jean-Yves Bouguet):** This segment of the course will describe practical techniques for acquiring accurate surface light field data of real life objects. First, I will present a system that captures complete 3D models of objects. I will address issues such as scanner calibration, image processing, sensitivity and multiple scans registration. I will then describe a technique for acquiring the appearance images and aligning them with the 3D model. Finally, I will present an alternative acquisition technique based on image silhouette extraction that combines the 3D shape computation with the appearance capture. Although presented in the context of surface light field data acquisition, this session will cover general topics of geometry computation, passive (camera) and active (projector) device calibration that are of great interest in numerous applications in machine vision and computer graphics.

3. **Hardware-Accelerated Rendering of Surface Light Fields (Radek Grzeszczuk):** This segment will show a method for efficient representation and interactive visualization of surface light fields. In particular, we propose to approximate the radiance data by partitioning it over elementary surface primitives and decomposing each part into a small set of lower-dimensional discrete functions. We also propose a hardware-accelerated method of rendering from this compact representation that accurately conveys the physical realism of the original data at interactive frame rates on a personal computer. Finally, we show that our representation can be further compressed using standard image compression techniques leading to extremely compact data sets that are up to four orders of magnitude smaller than the uncompressed light field data. We demonstrate the approximations for a variety of non-trivial synthetic scenes and physical objects scanned through 3D photography.

4. **Acquisition of Light Field Data using Hand-Held Camera (Marc Pollefeys):** This segment of the course will demonstrate how light field data can be acquired using a hand-held camera. The presented approach combines different state-of-the-art algorithms, mainly developed in the field of computer vision, to automatically retrieve all the necessary data from the raw images. The approach starts by relating neighboring views based on automatically extracted features. From this both the motion and the calibration of the camera are computed. In a second stage the surface of the observed scene is also estimated from the images using a multi-view stereo algorithm. The obtained results can be used to generate different types of visual models. In particular, we will propose an unstructured light field approach that uses view-dependent geometry approximations.

5. **Image-Based 3D Photography using Opacity Hulls**
Leonard McMillan, MIT, and Hanspeter Pfister, MERL

This segment of the course introduces a new image-based approach to capturing and modeling highly specular, fuzzy, transparent, or translucent objects. We present a system for automatically acquiring high quality graphical models of objects that are extremely difficult to scan with traditional 3D scanners. The system consists of turntables, a set of cameras and lights, and monitors to project colored backdrops. We use multi-background matting techniques to acquire alpha mattes and images of the object from multiple viewpoints. The alpha mattes are used to construct an *opacity hull*. The opacity hull is a new shape representation, defined as the visual hull of the object with view-dependent opacity. It enables visualization of complex object silhouettes and seamless blending of objects into new environments. The system also supports relighting of objects with arbitrary appearance using surface reflectance fields, a purely image-based appearance representation. This system is the first to acquire and render fuzzy, transparent, and translucent 3D objects, such as feathers or a glass of beer, from arbitrary viewpoints under

Course Schedule

Time	Topic	Speaker
1:30	Introduction	Grzeszczuk
1:30	Eigen-Texture Method for Compression and Synthesis of Reflectance Data	Nishino
2:15	Acquisition of Surface Light Fields	Bouguet
2:45	Hardware-Accelerated Rendering of Surface Light Fields	Grzeszczuk
3:15	Break	
3:30	Acquisition of Light Field Data using Hand-Held Camera	Pollefeys
4:15	Image-Based 3D Photography using Opacity Hulls (Part I)	McMillan
4:45	Image-Based 3D Photography using Opacity Hulls (Part II)	Pfister
5:15	Adjourn	