Proposal for a SIGGRAPH 2008 Course:
Computational Photography: Advanced Topics

Online ID: courses_0040

1 Course Title

Computational Photography

2 Category

Image-based and 3D Photography techniques - other

3 Course Organizer

Ramesh Raskar
MIT-Media Lab (preferred)
20 Ames St.,
Cambridge, MA 02139, USA

Mitsubishi Electric Research Labs (MERL)
201 Broadway
Cambridge, MA 02139, USA

Jack Tumblin
EECS Dept, 3-320 Ford Design Center
Northwestern University
2133 Sheridan Road
Evanston, IL 60208, USA

4 Proposed Length

3.75 Hour

5 Proposed Presentation Venue

Regular session room

6 Summary Statement

Participants learn about the latest computational methods in digital imaging that overcome the traditional limitations of a camera and enable novel imaging applications. The course provides a practical guide to topics in image capture and manipulation methods for generating compelling pictures for computer graphics and for extracting scene properties for computer vision, with several examples.

7 Names of Lecturers

- Paul DEBEVEC(University of Southern California, Institute for Creative Technologies (USC-ICT), USA)(paul@debevec.org)
- Ramesh RASKAR (Mitsubishi Electric Research Labs (MERL), USA) (raskar@merl.com)
- Jack TUMBLIN (Northwestern University, USA) (jet@cs.northwestern.edu)

8 Course Abstract

Abstract

Computational photography combines plentiful computing, digital sensors, modern optics, many varieties of actuators, probes and smart lights to escape the limitations of traditional film cameras and enables novel imaging applications. Unbounded dynamic range, variable focus, resolution, and depth of field, hints about shape, reflectance, and lighting, and new interactive forms of photos that are partly snapshots and partly videos, performance capture and interchangeably relighting real and virtual characters are just some of the new applications emerging in Computational Photography. The computational techniques encompass methods from modification of imaging parameters during capture to sophisticated reconstructions from indirect measurements.

For the first time, we will bypass basic and introductory material presented in earlier versions of this course (Computational Photography 2005,6,7) and expand coverage of more recent topics. Emphasizing more recent work in computational photography and related fields (2006 or later) this course will give more attention to advanced topics only briefly touched before, including tomography, heterodyning and Fourier Slice applications, inverse problems, gradient illumination, and novel optics. With this deeper coverage, the course offers a diverse but practical guide to topics in image capture and manipulation methods for generating compelling pictures for computer graphics and for extracting scene properties for computer vision, with several examples.

A larger problem is that a multi-disciplinary field that combines ideas from computational methods and modern digital photography involves a steep learning curve. While basic techniques are becoming well-known, many advanced ideas in computational photography are still relatively new to digital artists and programmers and there is no up-to-date reference text. For example, photographers are not always familiar with advanced algorithms now emerging. These topics, however, can be easily learned without extensive background. The goal of this presentation is to present both aspects in a compact form.

We introduce benefits of higher dimensional representation of light-fields and reflectance fields, clarify concepts such as ray-transfer matrix and explain wavefront coding and non-linear optics. We discuss applications of sensors for depth, thermal and millimeter waves. This is the first time an advanced course is being proposed. A subset of speakers presented a course in Computational Photography 2005-2007, where the emphasis was on learning about Siggraph/Computer Vision literature. In the proposed advanced course, our emphasis is on ideas and methods from outside the computer graphics and computer vision community.

The participants learn about topics in image capture and manipulation methods for generating compelling pictures for computer graphics and for extracting scene properties for computer vision, with several examples. We hope the course will better satisfy the technical specialist without intimidating the curious graphics researcher and any SIGGRAPH attendee interested in recent advances in photography. The intended audience is photographers, digital artists, image processing programmers and vision researchers using or building applications for digital cameras or images. They will learn about camera fundamentals and powerful computational tools, along with many real world examples.

9 Prerequisites

A basic understanding of camera operation and image processing is required.
Familiarity with concepts of linear systems, convolution, and machine vision will be useful.

10 Level of Difficulty

Intermediate (5 out of 10)

11 Intended Audience

Photographers, digital artists, image processing programmers and vision researchers using or building applications for digital cameras or images will learn about camera fundamentals and powerful computational tools, along with many real world examples.

12 Topics Covered

Camera internals
Camera parameters
Mathematical and image processing tools

Improving camera's dynamic range, framerate, resolution and focus
Introduction to applied optics
Image reconstruction techniques
Camera arrays and Multi-image fusion methods
Future smart and unconventional cameras

13 Course Syllabus

A.1 Introduction (Raskar, 15 minutes)

Digital photography compared to film photography
Image formation, Image sensors and Optics

A.2 Concepts in Computational Photography (Tumblin, 15 minutes)

The 'Photographic Signal'
What is the ideal photograph?
Ray-based versus pixel-based concepts
Understanding dimensionality of rays outside and inside the camera

A.3 Optics: Computable extensions (Raskar, 30 minutes)

Wavefront coding
Nonlinear optics; graded-indes,
'folded' optics; Tombo
Schlieren optics; extensions
Photonic Crystals and black silicon, negative index materials
Mirage program
Agile Spectrum Imaging
Random-lens Imaging (Torralba…)
What can we learn from animal eyes?

A.4 Sensor Innovations (Tumblin, 30 minutes)

trends in sensor pixel pitch (now ~1.9 micron)
resolution vs. noise issues;
assorted pixels and 'de-mosaicking',
hi-res 'streaming cameras'
Thermal sensors and benefits of thermal imaging;
Millimeter wave sensors, 3D sensors: Canesta, Zcam
Single-Photon Detectors
Penrose Tiles as pixels (Ben-Ezra),
Compressed Sensing

Q&A (5 minutes)
Break

B.1 Illumination as Computing (Debevec, 25 minutes)

Light stages,
Structured light for shape, reflectance, and more.
Coherent light advantages/problems;
femtosecond light sources & detection;
direct/indirect separations,
diffuse/specular separations,
Glare: Sensing, Compensation, and Control
True light-field displays (360' LF display),
pico-projectors; what is newly possible?

B.2 Scene & Performance Capture (Debevec 20 minutes)

gradient illumination methods
high-speed hybrids for real-time markerless performance capture;
visible and invisible markers: UV sensitive dye, etc.
tradeoffs: measurement time/cost/resolution/flexibility;

B 3. Image Aggregation & Sensible Extensions (Tumblin 20 minutes)

Merging uncalibrated cameras:
PhotoTourism,
knowledge from Web-cam aggregates (Robert Pleiss) ,
Generalized Dynamic Stitching methods for mismatched viewpoints,
Learning with databases (Alosha Efros papers: -pop-up 3D, re-lighting)
Face detection for auto-focus,
auto-exposure, Smile detection (Sony), Thinning Camera (HP) Seam Carving (MERL), etc.

B .4 Community and Social Impact (Raskar, 20 minutes)

CMU's captcha-like games for object recognition;
Google Earth problems; from street maps to street-level photos to 3D models
Is loss of privacy unavoidable? How can we compute a better result ?(Shai Avidan's work, etc.)
User inputs; applied 3D pose/motion capture, Cell-phone Mouse, Collision Avoidance devices
actively helpful illuminators: Contrast enhancing lamps (agile spectrum),

B.5. Summary and Discussion (All, 10 minutes)

new/next questions, new/next tools, future trends, advanced concepts, whats coming, whats next..

Topics not covered: film cameras, optical design, traditional image processing, image based rendering (IBR) and novel view synthesis, hardware technologies for lighting, projector-camera systems, geometric calibration and photometric calibration techniques, compression, storage.

14 Subtitles of Course Sections

A.1 Introduction (Raskar, 15 minutes)
A.2 Concepts in Computational Photography (Tumblin, 15 minutes)
A.3 Optics: Computable extensions (Raskar, 30 minutes)
A.4 Sensor Innovations (Tumblin, 30 minutes)

Q&A (5 minutes)
Break

B.1 Illumination as Computing (Debevec, 25 minutes)
B.2 Scene & Performance Capture (Debevec 20 minutes)

B 3. Image Aggregation & Sensible Extensions (Tumblin 20 minutes)
B .4 Community and Social Impact (Raskar, 20 minutes)
B.5. Summary and Discussion (All, 10 minutes)

15 Course Length Reduction

The course is being presented in a very compact format, and we hope the reviewers can keep the full length as it was reduced from last year's full-day course that received very good response . We have reduced the number of speakers from six to three to encourage focus on advanced topics. If the course length must be reduced to 2 hours, we can reduce the number and variety of example systems explored within each section.

16 Extant Materials

The course notes will include extended bibliography, psuedo-code and sample code snippets. An update will be made available on the course website. We hope the website will be useful for interested photographers, artists and programmers who are experimenting with creative camera uses and image processing applications in their own studios, universities, labs or companies. (Sample slides, bibliography and code is attached.)

17 Course History

The course follows well-received shorter half-day courses presented at Siggraph 2005 and 2006 (Attendee feedback is attached), and extended to a full-day course with 6 speakers in 2007. Despite a Sunday morning 8:30am schedule each year, the course was very well attended and the sessions were interactive. Several attendees asked us to consider a longer course with more details and deeper coverage. We submitted a Siggraph 2006 proposal for a full-day course after adding two prolific contributors to the field, Shree Nayar and Marc Levoy. We were heartened to see a positive review with scores of 4.0, 4.0, 4.0, 4.0, 4.0. However, the course was cut to a half-day course. For Siggraph 2007, our full-day course proposal added two speakers with very different backgrounds, and was accepted, with good audience response and strong attendance. Course evaluations suggested that more advanced material would be welcome, prompting us to invite a new speaker (Paul Debevec) with deep expertise.

There has always been interest in the Siggraph community for image-based applications. But, thanks to the growing prevalence of digital cameras, there has recently been a renewed interest in digital photography-based research and products. The papers at Siggraph conference include high dynamic range, matting, image fusion, synthetic aperture using camera arrays, flash photography and cartooning. A more detailed list is included in the sample bibliography. We plan to give an overview of these publications and papers at Computer Vision conferences, as well as topics in scientific imaging beyond photography. The speakers have published papers at Siggraph and other major venues in computational photography and are active in the research area.