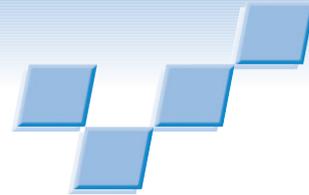


Computational Photography



Frédo Durand
Massachusetts Institute of Technology

Richard Szeliski
Microsoft Research

Since their inception over three decades ago, the fields of computer vision, image processing, and computer graphics have been concerned with analyzing, manipulating, and synthesizing images using numerical algorithms. While these three fields continue to evolve in a loosely coupled manner, users' near-ubiquitous access to digital cameras and personal computers has spurred renewed interest in what is now dubbed *computational photography*, which lies at the intersection of these fields.

Computational photography's precise definition is subject to as much discussion as the determination of its birth. One can argue that when Land and McCann turned their Retinex theory into an algorithm that reduced images' dynamic range, they pioneered intelligent image enhancement. When Burt and Adelson introduced pyramid-based image merging, they prefigured techniques such as Poisson image editing and digital photomontage, which create composite images achievable only through computational means.

Computational techniques that merge photographs to compensate for camera limitations have had a particular impact on users. Panorama stitching, derived from aerial photogrammetry techniques, was first applied to video by the Salient Stills project and Mann and Picard, and it culminated in the completely automatic recognition and creation of panoramas that Brown and Lowe demonstrated. Digitally combining multiple exposures enabled capturing high-dynamic-range images—as Mann and Debevec and Malik developed—while tone mapping algorithms let casual users create pictures that would require much skill and lighting equipment using traditional photography.

Traditional image processing topics, such as resolution enhancement and denoising, are receiving renewed interest, in particular using machine learning. Camera shake and motion blur are two critical issues in image quality; new deblurring algorithms can remove image blur based on models of natural image statistics, asergus and colleagues recently demonstrated. Other techniques that compensate for image blur include the coded shutter by Raskar and colleagues, and Ben-Ezra and Nayar's combining different spatial and temporal resolutions.

Separating foreground and background layers, also known as matting, is an issue as old as special effects,

but it also has many computational photography applications. Matting has received a wealth of new treatment in recent years based on local image statistics (as Ruzon and Tomasi and Chuang and colleagues proposed) and nonhomogeneous optimization (as Levin and colleagues demonstrated). Similarly, an image region's precise selection is now greatly simplified using techniques such as Mortensen and Barrett's intelligent scissors and Rother and colleagues' GrabCut.

Another major trend in computational photography and computational imaging in general has been the design of special optics and sensors that computers are meant to process before display. These methods include the catadioptric (mirror/lens) and generalized mosaicing systems developed in particular in Nayar's group. Cathey and Dowski's wavefront coding greatly increases field depth using modified optics and deconvolution. Pioneered by Wang and Adelson and refined by Levoy and colleagues, light field capture using lenslet arrays or multiple-camera systems enables 3D re-rendering as well as refocusing. Researchers such as Debevec and his team have shown that a user can capture scenes or objects under multiple illuminations and use them to perform relighting after the capture. A much simpler version of this idea is to simply combine a pair of images taken with and without a flash exposure, as Petschnigg and colleagues and Eisemann and Durand showed. Specialized optics and lighting have long been used in scientific fields such as astronomy, synthetic aperture radar, and microscopy. Ideas from these fields are slowly beginning to percolate into computational photography as well.

One of the most exciting recent developments in computational photography has been the gradual migration of computation algorithms from computers to cameras. For example, some cameras today use face detection to better focus and expose the image, while others perform preliminary panorama stitching directly in the camera and use local tone mapping to manage difficult lighting situations.

In this issue

This special issue includes three articles describing recent trends in computational photography.

In "Editing Soft Shadows in a Digital Photograph," Mohan and Tumblin present a technique to manipulate

Additional Resources

The Internet offers a wealth of resources on computational photography. Many courses post lecture notes online.

Courses

- Massachusetts Institute of Technology, Digital and Computational Photography: <http://groups.csail.mit.edu/graphics/classes/CompPhoto/>
- Carnegie Mellon University, Computational Photography: <http://graphics.cs.cmu.edu/courses/15-463/>
- University of Washington, Computer Vision: <http://www.cs.washington.edu/education/courses/cse576/>
- Georgia Institute of Technology, Computational Photography: http://www-static.cc.gatech.edu/classes/AY2006/cs4475_summer/
- University of Toronto, Introduction to Visual Computing: <http://www.cs.toronto.edu/~kyros/courses/320/>

Other material

- Raskar and Tumblin's computational photography primer: <http://www.merl.com/people/raskar/photo/>
- Slides from the 2005 Symposium on Computational Photography: <http://photo.csail.mit.edu/>

Tumblin and Raskar also have recently published a state-of-the-art report that covers all aspects of computational photography. *Computer* magazine (<http://www.computer.org/computer>) published a special issue on computational photography in August 2006 with articles summarizing work on time multiplexing for relighting by Debevec, the moment camera by Cohen and Szeliski, light field imaging by Levoy, and computational cameras by Nayar.

shadow boundaries in images. Their system lets the user loosely mark shadow boundaries and uses nonlinear optimization to finely model soft shadows' intensity falloff. By working in the image gradient domain, they can remove or modify shadows and preserve the fine detail of the underlying surfaces.

In "Optical Splitting Trees for High-Precision Monocular Imaging," McGuire et al. address the design of complex imaging systems in which multiple beam splitters

enable capturing a scene through the same optical axis but with different imaging parameters. The authors use optimization and take into account the real characteristics of the optical elements to devise an optimal configuration. They demonstrate applications such as high-dynamic-range imaging, focusing, matting, and high-speed imaging.

In "Exploring Defocus Matting: Nonparametric Acceleration, Super-Resolution, and Off-Center Matting," Joshi et al. present techniques to extract high-quality mattes using defocus information derived from multiple coaxial video streams. They exploit temporal coherence to greatly accelerate the matte computation by building a nonparametric model from a small number of key frames. This technique also enables the use of cameras with different resolutions so that only one camera must be high resolution, and it makes possible the combining of a coaxial defocus matting camera with one off-centered master camera. ■



Frédo Durand is an associate professor in the Electrical Engineering and Computer Science Department at the Massachusetts Institute of Technology, where he is a member of the Computer Science and Artificial Intelligence Laboratory. His research interests include realistic graphics, real-time rendering, nonphotorealistic rendering, and computational photography. Durand has a PhD from Grenoble University, France. Contact him at fredo@mit.edu.



Richard Szeliski is a principal researcher and leads the Interactive Visual Media Group at Microsoft Research. His research interests include digital and computational photography, video scene analysis, 3D computer vision, and image-based rendering. Szeliski has a PhD in computer science from Carnegie Mellon University. Contact him at szeliski@microsoft.com.



REACH HIGHER

Advancing in the IEEE Computer Society can elevate your standing in the profession.

Application to Senior-grade membership recognizes

- ✓ ten years or more of professional expertise

Nomination to Fellow-grade membership recognizes

- ✓ exemplary accomplishments in computer engineering

GIVE YOUR CAREER A BOOST ■ UPGRADE YOUR MEMBERSHIP

www.computer.org/join/grades.htm