

# Computational Photography Comes into Focus

*Advances in computational photography are making image capture the starting point. The technology is transforming the field.*

**O**VER THE LAST decade, digital cameras have radically refocused the way people capture and manipulate pictures. Today, the snap of a photo is merely a starting point for composing and manipulating an image. A photographer can make basic changes to a picture from within the camera, but also may use photoediting software on a computer to significantly alter the look, feel and composition. “We can use computation to make the process better, both aesthetically and in terms of greater flexibility,” explains Frédo Durand, a professor in the Computer Science and Artificial Intelligence Laboratory at MIT in Cambridge, MA.

Researchers and engineers are now taking the concept further. They are designing different types of cameras, developing increasingly sophisticated algorithms, and using new types of sensors and systems to boldly go where no camera has gone before. The ability to record richer information about a scene and use powerful image enhancement techniques are redefining the field. “Computational photography and computational imaging are extremely vibrant areas,” states Shree K. Nayar, professor of computer science at Columbia University in New York City.

These cameras, along with more advanced software, will radically change the way people view and use images. For example, they will make it possible to detect a tiny object or imperceptible motion from the field of view. They might change the perspective or angle after a photo is snapped, or provide a 360-degree panoramic view. They might also augment reality and refocus various objects in scenes, after a photo has been shot.



The Lytro camera captures the entire light field.

Meanwhile, smartphone cameras will further redefine photography by incorporating sensors and greater onboard computational power. Combined with specialized apps or cloud-based services, they will stretch the current concept of photography in new and intriguing ways.

## A Better Image

It is no secret that digital cameras have reinvented photography. The transition from film to pixels has created an opportunity to manipulate and share photos in ways that were not imaginable in the past. However, today's cameras rely heavily on the same features and image capture techniques as film cameras; they are largely designed the same way film cameras were, but with new features. “They present a lot of limitations. It is very difficult to change

the way the camera behaves or the way it captures images,” Durand explains.

However, the use of computational photography, imaging, and optics promises to significantly change the way people approach photography, capture images, and edit them. For example, William Freeman, a professor of computer science at MIT, says computational cameras could capture multiple images at a time to compensate for glare, oversaturation, and other exposure problems. They could also eliminate the need for a flash. “Too often, flash ruins the tonal scale of images,” he says, “but by combining multiple shots, both with flash and without, it is possible to create a single sharp, low-noise image that has a beautiful tone scale.”

Similarly, the ability to change focus after capturing a shot would make it

possible to fix on a person in the foreground while also focusing on an object in the distance, like the Eiffel Tower or Statue of Liberty; everything else in the photo would appear blurred. The commercially available Lytro camera—which records the entire light field in the frame (essentially, depth of field data about the entire scene)—already allows a user to refocus pictures and adjust lighting after image capture. Likewise, a sensor that would capture different levels of light on different pixels could create entirely new types of photographs, including images with markedly different brightness and color ranges.

The technology of computational photography could also lead to changes in camera design. As Columbia's Nayar points out, computational features alone deliver significant improvements, but they also create the possibility for new types of camera bodies, lenses, and optics. Adding a computational lens to a smartphone, for instance, could mimic the high-end features of an expensive optical lens at a much lower price point, or may create entirely new features. A photographer might snap on a lens or multiple lenses that would provide 3-D capabilities, or marry video and still photography to address camera shake, particularly in difficult low-light or high-speed environments.

The benefits of computational cameras and software are likely to extend far beyond consumers. The technology could impact an array of industries, including medicine, manufacturing, transportation and security, points out Marc Levoy, a professor of computer science and electrical engineering at Stanford University in Palo Alto, CA, who recently took leave to work with the Google Glass development team. Levoy says cameras with more advanced computational capabilities could redefine the way we think about the world around us, and provide insights that extend beyond basic images or video.

For example, he and other researchers have explored the idea of developing a computational camera that could see through crowds, objects, and people. The technology could also generate a focal stack within a single snapshot. This could create new opportunities

## Computational photography could lead to changes in camera design, such as new types of camera bodies, lenses, and optics.

in biology and microscopy, Levoy says. "A technician could capture images of cell cultures without focusing a microscope; focusing would take place after the picture is taken." A computational camera could also automatically count the number of cells in an image and provide information faster and more accurately than any human, he adds.

Perhaps the highest-profile example of a computational photography system to date is Google Glass. Its camera captures images and provides additional information and insight in an array of situations and scenarios—a step toward more-advanced augmented reality tools. Among other things, the Google Glass team is focused on developing map data, language translations, travel and transit information,

and apps that track health, exercise data and body information. The device also can capture a burst of images and deliver improved high-dynamic-range imaging and low-light imaging.

### Beyond Pixels

Engineering these systems and developing the algorithms to support these devices is no simple task, particularly as researchers look to extend computational capabilities beyond the world of consumer cameras into fields such as astronomy, medical photography, and automobile photography. There also is the possibility of capturing images beyond the visible spectrum of light, incorporating environmental sensors, or finding ways to apply algorithms to detect small but important changes in the environment. As Levoy puts it, "There is a potential for this technology to be extremely disruptive."

Durand also says the gains are not limited to conventional cameras. New types of cameras and software could generate robust 3-D images that reveal things not visible through optics alone. Already, he and Freeman have developed algorithms that can sense the flow of blood in a person's face, or detect one's heartbeat based on subtle head motions. This relates to a technique called *motion magnification*, which could potentially be used to detect weaknesses in bridges and build-



A cutaway view of the Canon EOS 5D Mark II camera body.

IMAGE COURTESY OF CANON USA

ings; it amplifies pulse signals and color variations. “These signals cannot be detected by the human eye, but they are revealed through advanced computational imaging and slow-motion analysis,” Freeman explains.

Vladimir Katkovnik, a professor of signal processing at Tampere University of Technology in Finland, says a significant hurdle to accomplishing all this is the development of algorithms that sort through all the data and apply it in usable ways. Despite the prospect of larger sensors that can capture more data, there is a trend toward more pixels in images. “Larger numbers of megapixels means images with more pixels of a smaller size. As smaller numbers of photons appear on a pixel during exposure time, there is a larger amount of noise generated. Noise removal is a growing challenge in any imaging or sensing device; the end quality depends on how well noise is removed.”

Another challenge, Durand says, is developing robust algorithms that work effectively on relatively small devices such as cameras, smartphones, and tablets. “The issue is not necessarily whether you can develop an algorithm that works; it is whether it is possible to map the computation to the hardware in an efficient manner. Writing optimized code that can take advantage of modern hardware, including mobile processors, is extremely difficult.” He is currently developing a compiler to make it easier to achieve high performance, without devoting a large development team to the task.

Nayar believes researchers will tap into big data techniques and, in some cases, examine and analyze existing photos to build algorithms that drive even more sophisticated image processing. Right now, “if you try to remove a person or object from a photo, there is no easy way to fill the hole, even with fairly sophisticated photoediting software,” he says. “By using millions of pictures and applying machine learning algorithms, it is possible to fill the holes in visually plausible ways.” At some point, he adds, these capabilities will likely appear on cameras, smartphones, and tablets, and provide nearly instantaneous manipulation and editing tools that make today’s image-editing options pale by comparison.

Researchers are likely to hit the tipping point within the next decade, as increasingly powerful processors and a greater knowledge of physics push the technology forward. “The algorithms being used today are still mostly in the infant stages,” Nayar says. “So far, most of the research has revolved around extending the capabilities of traditional imaging and finding ways to improve the performance of digital cameras.” As knowledge about non-traditional imaging and optics converge, he notes, “everything from chip design to lens and camera design will undergo major changes.”

In the end, Durand says it is important to place computational photography, imaging, and optics in the right context. The technology will not replace today’s cameras and photographs; it will enhance them and continue advancing a process that dates back thousands of years, to the development of pinhole cameras. Computational photography puts data to use in new and better ways, whether it is applied to DNA sequencing or to improved traffic cameras or security tools.

Says Durand, “Photography is just one aspect of a much bigger picture. With it, we are able to see the world in a fundamentally different way.” **C**

#### Further Reading

Ragan-Kelley, J., Adams, A., Paris, S., Levoy, M., Amarasinghe, S., Durand, F.

Decoupling Algorithms from Schedules for Easy Optimization of Image Processing Pipelines, SIGGRAPH 2012, <http://people.csail.mit.edu/jrk/halide12/halide12.pdf>.

Bychkovsky, V., Paris, S., Chan, E., Durand, F. When Does Computational Imaging Improve Performance?, IEEE Transactions on Image Processing, 2012. [http://www1.cs.columbia.edu/CAVE/publications/pdfs/Cossairt\\_TIP12.pdf](http://www1.cs.columbia.edu/CAVE/publications/pdfs/Cossairt_TIP12.pdf)

Cossairt, O., Gupta, M., Nayar, S.K. Ironies of automation. *New Technology and Human Error*, J. Rasmussen, K. Duncan, J. Leplat (Eds.). Wiley, Chichester, U.K., 1987, 271–283.

Cho, T.S., Avidan, S., Freeman, W.T. A Probabilistic Image Jigsaw Puzzle Solver, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2010. <http://people.csail.mit.edu/billf/papers/JigsawSolverCVPR2010.pdf>

Samuel Greengard is an author and journalist based in West Linn, OR.

© 2014 ACM 0001-0782/14/02 \$15.00

#### Opportunity

## Heidelberg Laureate Forum

Individuals may apply through the end of February for one of 200 openings to attend the second Heidelberg Laureate Forum (HLF), to be held Sept. 21–26 at Heidelberg University in Germany.

HLF allows researchers from all over the world to interact with laureates of the most prestigious awards in computer science and mathematics. Last year, 40 laureates, including recipients of the ACM A.M. Turing Award, the International Mathematical Union’s Fields Medal and Nevanlinna Prize, the Norwegian Academy of Science and Letter’s Abel Prize, addressed young researchers on topics that ranged from “how to do research” to deep technical areas of science and math.

ACM Europe chair Fabrizio Gagliardi, one of a number of ACM members (among others) who helped organize the initial Forum, said about last year’s event, “I was impressed by the attitude of the laureates who all spent a considerable amount of time networking with the young researchers; not only during the HLF sessions, but also during meals and in after-dinner discussions. This is probably the highest value of the event: providing a relative small set of promising future scientists with the unique opportunity to engage with some of the most brilliant minds in mathematics and computer science.”

Interested researchers may apply on the HLF website, at <https://application.heidelberg-laureate-forum.org>.

HLF is organized by the Heidelberg Laureate Forum Foundation in cooperation with Klaus Tschira Stiftung and the Heidelberg Institute for Theoretical Studies, as well as ACM, the International Mathematical Union, and The Norwegian Academy of Science and Letters.

—Lawrence Fisher