# N news

Kirk L. Kroeker

Photography's Bright Future

Researchers working in computational photography are using computer vision, computer graphics, and applied optics to bring a vast array of new capabilities to digital cameras.

WHILE THE FIRST digital cameras were bulky, possessed few features, and captured images that were of dubious quality, digital cameras today pack an enormous array of features and technologies into small designs. And the quality of pictures captured on these cameras has improved dramatically, even to the point where most professional photographers have abandoned film and shoot exclusively with digital equipment. Given that digital photography has established itself as superior to analog film in many aspects, it might seem safe to assume that the next breakthroughs in this area will be along the lines of more megapixels or smaller handheld designs. However, researchers working in the emerging area of computational photography—a movement that draws on computer vision, computer graphics, and applied optics—say the next major breakthroughs in digital photography will be in how images are captured and processed.

Indeed, the technology powering digital photography is rapidly improving and is certainly facilitating the ability to capture images at increasingly
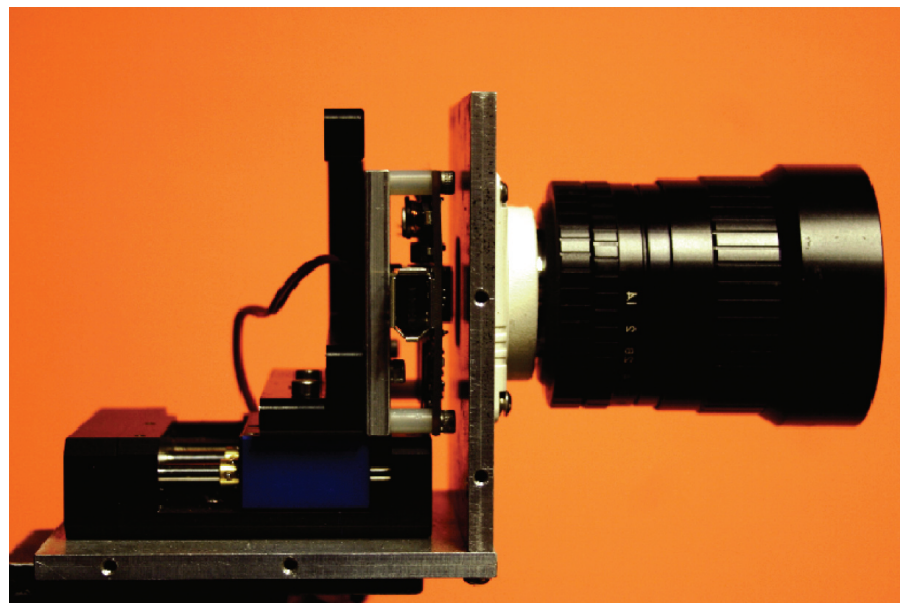
A camera mechanism designed to enhance depth of field without compromising light quality. A sensor mounted on a movable platform and controlled by a microactuator can capture an image that is equally blurred everywhere but can be deblurred to produce an image with an unusually large depth of field.

high resolutions on ever smaller hardware. But most digital cameras today still operate much like traditional film cameras, offering a similar set of features and options. Researchers working in computational photography are pushing for new technologies and de-

signs that will give digital cameras abilities that analog cameras do not have, such as the ability to capture multiple views in one image or change focal settings even after a shot is taken.

"There is tremendous enthusiasm for computational photography," says

Shree Nayar, a professor of computer science at Columbia University. "The game becomes interesting when you think about optics and computation within the same framework, designing optics for the computations and designing computations to support a certain type of optics." That line of thinking, Nayar says, has been evolving in both fields, optics on one side and computer vision and graphics on the other.

Multiple Images

One of the most visually striking examples of computational photography is high dynamic range (HDR) imaging, a technique that involves using photo-editing software to stitch together multiple images taken at different exposures. HDR images—many of which have a vibrant, surreal quality—are almost certain to elicit some degree of surprise in those who aren't yet familiar with the technique. But the traditional process of creating a single HDR image is painstaking and cannot be accomplished with moving scenes because of the requirement to take multiple images at different exposures. In addition to being inconvenient, traditional HDR techniques require expertise beyond the ability or interest of casual photographers unfamiliar with photo-editing software. However, those working in computational photography have been looking for innovative ways not only to eliminate the time it takes to create an HDR image, but also to sidestep the learning curve associated with the technique.

With computational photography, people can change a camera's focal settings after a photo is taken.

"It turns out that you can do HDR with a single picture," says Nayar. "Instead of all pixels having equal sensitivity on your detector, imagine that neighboring pixels have different sunshades on them—one is completely open, one is a little bit dark, one even darker, and so on." With this technique, early variations of which have begun to appear in digital cameras, such as recent models in the Fujifilm FinePix line, the multiple exposures required of an HDR image would be a seamless operation initiated by the user with a single button press.

Another research area in computational photography is depth of field. In traditional photography, if you want a large depth of field—where everything in a scene is in focus—the only way to do so is to make the camera's aperture very small, which prevents the camera from gathering light and causes images to look grainy. Conversely, if you want a good picture in terms of brightness and color, then you must open the camera's aperture, which results in a reduced depth of field. Nayar, whose work involves developing vision sensors and creating algorithms for scene interpretation, has been able to extend depth of field without compromising light by moving an image sensor along a camera's optical axis. "Essentially what you are doing is that while the image is being captured, you are sweeping the focus plane through the scene," he explains. "And what you end up with is, of course, a picture that is blurred, but is equally blurred everywhere." Applying a deconvolution algorithm to the blurred image can recover a picture that Nayar says doesn't compromise the quality of the image.

One of the major issues that those working in computational photography face is testing their developments on real-world cameras. With few exceptions, the majority of researchers working in this area generally don't take apart cameras or try to make their own, which means most research teams are limited to what they can do with existing cameras and a sequence of images. "It would be nicer if they could program the camera," says Marc Levoy, a professor of computer science and electrical engineering at Stanford University. Levoy, whose research involves light-field sensing and applications of computer graphics in microscopy and biology, says that even those researchers who take apart cameras to build their own typically do not program them in real time to do on-the-spot changes for different autofocus algorithms or different metering algorithms, for example.

"No researchers have really addressed those kind of things because they don't have a camera they can play with," he says. "The goal of our open-source camera project is to open a new angle in computational photography by providing researchers and students with a camera they can program." Of course, cameras do have computer software in them now, but the vast majority of the software is not available to researchers. "It's a highly competitive, IP-protected, insular industry," says Levoy. "And we'd like to open it up."

But in developing a programmable platform for researchers in computational photography, Levoy faces several

A photo containing 62 frames, taken through a group of trees, of Chicago's Newberry Library.

major challenges. He says the biggest problem is trying to compile code written in an easy-to-use, high-level language down to the hardware of a camera. He likens the challenge to the early days of shading languages and graphics chips. Shaders used to be fixed functions that programmers could manipulate in only certain limited ways. However, graphics chipmakers changed their hardware to accommodate new programming languages. As a result, says Levoy, the shader languages got better in a virtuous cycle that resulted in several languages that can now be used to control the hardware of a graphics chip at a very fine scale.

"We'd like to do the same thing with cameras," Levoy says. "We'd like to allow a researcher to program a camera in a high-level language, like C++ or Python. I think we should have something in a year or two."

In addition to developing an open source platform for computational photography, Levoy is working on applying some of his research in this area to microscopy. One of his projects embeds a microlens array in a microscope with the goal of being able to refocus photographs after they are taken. Because of the presence of multiple microlenses, the technology allows for slightly shifting the viewpoint to see around the sides of objects—even after capturing an image. With a traditional microscope, you can of course refocus on an object over time by moving the microscope's stage up and down. But if you are trying to capture an object that is moving or a microscopic event that is happening very quickly, being able to take only a single shot before the scene changes is a serious drawback. In addition to Levoy's microlens array allowing images to be refocused after they are captured, the technology also offers the ability to render objects in three dimensions. "Because I can take a single snapshot and refocus up and down," he says, "I can get three-dimensional information from a single instant in time."

These and other developments in computational photography are leading to a vast array of new options for researchers, industry, and photography enthusiasts. But as with any advanced technologies that have interfaces designed to be used by humans, one of

Marc Levoy's open-source camera project might enable researchers to program a camera in a high-level language, like C++ or Python.

the major challenges for computational photography is usability. While computer chips can do the heavy lifting for many of these new developments, the perception that users must work with multiple images or limitless settings to generate a good photo might be a difficult barrier to overcome. Levoy points to the Casio EX-F1 camera as a positive step toward solving this usability problem. He says the EX-F1, which he calls the first computational camera, is a game changer. "With this camera, you can take a picture in a dark room and it will take a burst of photos, then align and merge them together to produce a single photograph that is not as noisy as it would be if you took a photograph without flash in a dark room," he says. "There is relatively little extra load on the person."

Levoy predicts that cameras will follow the path of mobile phones, which, for some people, have obviated the need for a computer. "There are going to be a lot of people with digital cameras who don't have a computer and never will," he says. "Addressing that community is going to be interesting." He also predicts that high-end cameras will have amazing flexibility and point-and-shoot cameras will take much better pictures than they do now. Nayar is of a similar opinion. "One would ultimately try to develop a camera where you can press a button, take a picture, and do any number of things with it," he says, "almost like you're giving the optics new life." **C**

Kirk L. Kroeker works in communications and has written extensively about the impact of emerging technologies.

Cognitive Computing
IBM's Brain-Like Computer

IBM has received a $4.9 million grant from DARPA to lead an ambitious, cross-disciplinary research project to create a new computing platform: electronic circuits that operate like a brain.

"The mind has an amazing ability to integrate ambiguous information across the senses, and it can effortlessly create the categories of time, space, object, and interrelationship from the sensory data," Dharmendra Modha, the IBM researcher who is leading the project, told BBC News. "The key idea of cognitive computing is to engineer mind-like intelligent machines by reverse engineering the structure, dynamics, function, and behavior of the brain."

The cognitive computing initiative will include neuroscientists, computer and materials scientists, and psychologists from IBM and five U.S. universities. Modha believes the time is right for the project as neuroscientists have learned a great deal about the inner workings of neurons and their connecting synapses, resulting in "wiring diagrams" for the brains of simple animals. Also, supercomputing is able to simulate brains up to the complexity level of small mammals; last year a Modha-led IBM team used its Blue Gene supercomputer to simulate the 55 million neurons and half-a-trillion synapses in a mouse's brain. "But the real challenge is then to manifest what will be learned from future simulations into real electronic devices—nanotechnology," according to Modha.

Along with IBM Almaden Research Center and IBM T. J. Watson Research Center, Stanford University, University of Wisconsin-Madison, Cornell University, Columbia University Medical Center, and University of California-Merced are participating in the project.

Modha acknowledges that the project is very ambitious in terms of its scope and goals. "We are going not just for a home run," he said, "but for a home run with the bases loaded."