MVA2007 IAPR Conference on Machine Vision Applications, May 16-18, 2007, Tokyo, JAPAN

Computational Cameras

Shree K. Nayar
Department of Computer Science
Columbia University
New York, N.Y. 10027
Email: nayar@cs.columbia.edu

Abstract

The traditional camera is based on the principle of the camera obscura and produces linear perspective images. A computational camera uses unconventional optics to capture a coded image and software to decode the captured image to produce new forms of visual information. We show examples of computational cameras that capture wide field of view images, high dynamic range images, multispectral images, and depth images. We also describe through examples how the capability of a computational camera can be enhanced by using a controllable optical system for forming the image and a programmable light source as the camera's flash.

1. The Traditional Camera

Most cameras in use today are based on the principle of the camera obscura, which in Latin means "dark room." The concept of the camera obscura was first explored by Chinese philosophers in the 5th century B.C. and later by Arabian scientist-philosophers in the 11th century. It was only in the 16th century that it became known in the West, where it was turned into a powerful tool by artists to produce geometrically precise renditions of the real world [12]. In its earliest versions, the camera obscura was realized by piercing a pinhole in a wall to create a linear perspective image of the scene on a second wall. The artist could then walk up to the second wall and sketch out the image of the scene. While the camera obscura produced a clear image, it was a very dim one, as a pinhole severely limits the light energy that can pass through it. Within a matter of decades, the camera obscura was enhanced with the addition of a lens, which could collect more light and hence make the images brighter.

Over the next few centuries, the camera obscura went through many refinements that were geared towards making it easier for an artist to use. It is important to note that, in all of this, the artist was an essential part of the process of creating an image. From this viewpoint, the invention of film in the 1830s was a breakthrough. One could place a sheet of film exactly where the camera obscura formed an image of the scene and instantly record the image. That is, the artist was no longer an essential part of the process. This was clearly a very important moment in history. The advent of film made it remarkably easy to produce visual information and hence profoundly impacted our ability to communicate with each other and express ourselves.

It is often said that the invention of film was the most important event in the history of imaging. However, a few decades from now we may realize that a more significant invention took place around 1970 – the solid-state image detector. This device does exactly what film can do, except that one need not replace it each time a picture is taken. A single solid-state image detector can produce any number of images, without the need to develop or process each one. It took about 25 years for the image detector to mature into a reliable and cost-effective technology. Ultimately, in the mid 1990s, we witnessed an explosion in the marketplace of digital cameras. Today, one can go out and buy a digital camera for a few hundred dollars that fits into a shirt pocket and produces images that are comparable in quality to film.

2. Computational Cameras

We can all agree that, over the last century, the evolution of the camera has been truly remarkable. However, it is interesting to note that throughout this journey the principle underlying the camera has remained the same, namely, the camera obscura. As shown in Figure 1(a), the traditional camera has a detector (which could be film or solid-state) and a lens which only captures those principal rays that pass through its center of projection, or effective pinhole. In other words, the traditional camera performs a very special and restrictive sampling of the complete set of rays, or the light field [4], that resides in any real scene.

If we could configure cameras that sample the light field in radically different ways, perhaps, new and useful forms of visual information can be created. This brings us to the

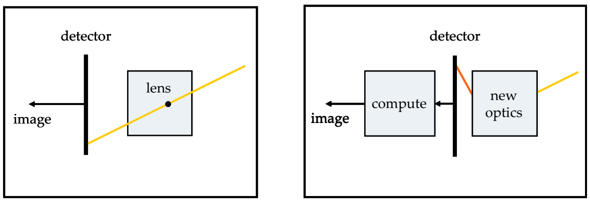

(a) The traditional camera. (b) A computational camera.

Figure 1. (a) The traditional camera is based on the principle of the camera obscura and produces a linear perspective image. (b) A computational camera uses novel optics to capture a coded image and a computational module to decode the captured image to produce new types of visual information.

notion of a computational camera [11], which is illustrated in Figure 1(b). It embodies the convergence of the camera and the computer. It uses new optics to map rays in the light field to pixels on the detector in some unconventional fashion. For instance, the yellow ray shown in the figure, which would have traveled straight through to the detector in the case of a traditional camera, is assigned to a different pixel. In addition, the brightness and spectrum of the ray could be altered before it is received by the pixel, as illustrated by the change in its color from yellow to red.

In all cases, the captured image is optically coded and hence, in its raw form, may not be easy to interpret. However, the computational module knows everything it needs to about the optics. Hence, it can decode the captured image to produce new types of images that could benefit a vision system. The vision system could be a human observing the images or a computer vision system that analyzes the images to interpret the scene.

In this article, I present a few examples of computational cameras that have been developed in collaboration with students and research scientists at the Computer Vision Laboratory at Columbia University. Imaging can be viewed as having several dimensions, including, spatial resolution, temporal resolution, spectral resolution, field of view, dynamic range and depth. Each of the cameras I present here can be viewed as exploring a specific one of these dimensions.

The first imaging dimension we will look at is field of view. Most imaging systems, biological as well as artificial ones, are rather limited in their fields of view. They can capture only a small fraction of the complete sphere around their location in space. Clearly, if a camera could capture the complete sphere or even a hemisphere, it would profoundly impact the capability of the vision system that uses it[1].

Related Work

There are several academic and industrial research teams around the world that are developing a variety of computational cameras. In addition, there are well established imaging techniques that naturally fall within the definition of a computational camera. A few examples are integral imaging [7] for capturing the 4D light field of a scene[2]; coded aperture imaging [2] for enhancing the signal-to-noise ratio of an image; and wavefront coded imaging [1] for increasing the depth of field of an imaging system. In each of these cases, unconventional optics is used to capture a coded image of the scene, which is then computationally decoded to produce the final image. This approach is also used for medical and biological imaging, where it is referred to as computational imaging. Finally, significant technological advances are also being made with respect to image detectors. In particular, several research teams are developing detectors that can perform image sensing as well as early visual processing (see [9][13][6] for some of the early work in this area).

When one thinks about wide-angle imaging, the fish eye lens [10] first comes to mind as it has been around for about a century. It uses what are called meniscus lenses to severely bend light rays into the camera, in particular, the rays that are in the periphery of the field of view. The limitation of the fish eye lens is that it is difficult to design one with a field a view that is much larger than a hemisphere while maintaining high image quality. The approach we have used is called catadioptrics. Catoptrics is the use of mirrors and dioptrics is the use of lenses. Catadioptrics is the combined use of lenses and mirrors. This approach has been been extensively used to develop telescopes [8]. While in the case of a telescope one is interested in capturing a very small field of view, here we are interested in exactly the opposite – the capture of an unusually large field of view.

In developing a wide-angle imaging system, it is highly desirable to ensure that the principal rays of light captured by the camera pass through a single viewpoint, or center of projection. If this condition is met, irrespective of how distorted the captured image is, one can use software to map any part of it to a normal perspective image. For that matter, the user can emulate a rotating camera to freely explore the captured field of view. In our work, we have derived a complete class of mirror-lens combinations that capture wide-angle images while satisfying the single viewpoint constraint. This family of cameras include ones that use ellipsoidal, hyperboloidal and paraboloidal mirrors, some of

[1]The French philosopher Michel Foucault has explored at great length the psychological implications of being able to see everything at once in his discussion of the panopticon [3].

[2]For recent advances in this approach, please see the article by Marc Levoy in this issue.

which were implemented in the past. We have also shown how two mirrors can be used to reduce the packaging of the imaging system while maintaining a single viewpoint.

A member of this class of wide-angle catadioptric cameras is shown on the left of Figure 2(a). It is implemented as an attachment to a conventional camera with a lens, where the attachment includes a relay lens and a paraboloidal mirror. As can be seen from the figure, the field of view of this camera is significantly greater than a hemisphere. It has a 220 degree field of view in the vertical plane and a 360 degree field of view in the horizontal one. An image captured by the camera is shown in the middle. The black spot in the center is the blindspot of the camera where the mirror sees the relay lens. Although the image was captured from close to ground level, one can see the sky above the bleachers of the football stadium. This image illustrates the power of a single-shot wide-angle camera over traditional methods that stitch a sequence of images taken by rotating a camera to obtain a wide-angle mosaic. While mosaicing methods require the scene to be static during the capture process, a single-shot camera can capture a wide view of even a highly dynamic scene.

Since the computational module of the camera knows the optical compression of the field of view achieved by the catadioptric system, it can map any part of the captured image to a perspective image, such as the one shown on the right of Figure 2(a). This mapping is a simple operation that can be done at video-rate using even a low-end computer. We have demonstrated the use of 360 degree cameras for video conferencing and video surveillance.

Another imaging dimension that is of great importance is dynamic range. While digital cameras have improved by leaps and bounds with respect to spatial resolution, they remain limited in terms of the number of discrete brightness values they can measure. Consider a scene that includes a person indoors lit by room lamps and standing next to an open window in which the scene outdoors is brightly lit by the sun. If one increases the exposure time of the camera to ensure the person appears well lit in the image, the scene outside the window would be washed out, or saturated. Conversely, if the exposure time is lowered to capture the bright outdoors, the person will appear dark in the image. This is because digital cameras typically measure 256 levels (8 bits) of brightness in each color channel, which is simply not enough to capture the rich brightness variations in most real scenes.

A popular way to increase the dynamic range of a camera is to capture many images of the scene using different exposures and then use software to combine the best parts of the differently exposed images. Unfortunately, this method requires the scene to be more or less static as there is no reliable way to combine the different images if they include fast moving objects. Ideally, we would like to have the benefits of combining multiple exposures of a scene, but with the capture of a single image.

In a conventional camera, all pixels on the image detector are made equally sensitive to light. Our solution is to create pixels with different sensitivities either by placing an optical mask with cells of different transmittances on the detector or by having interspersed sets of pixels on the detector exposed to the scene over different integration times. We refer to such a detector as one having an assortment of pixels. Note that most color cameras already come with an assortment of pixels – neighboring pixels have different color filters attached to them. In our case, the assortment is more complex as a small neighborhood of pixels will not only be sensitive to different colors but the pixels of the same color will have different transmittances or integration times as well.

A camera with assorted pixels is shown on the left of Figure 2(b). Unlike a conventional camera, in this case, for every pixel that is saturated or too dark there will likely be a neighboring pixel that is not. Hence, even though the captured image may have bad data, they are interspersed with the good data. An image captured with this camera is shown in the middle of the figure. In the magnified inset image one can see the expected checkerboard appearance of the image. By applying an image reconstruction software to this optically coded image a wide dynamic range image can be obtained, as shown on the right of the figure. Notice how this image includes details on the dark walls lit by indoor lighting as well as the bright sunlit regions outside the door.

Figure 2(c) shows how the well known method of image mosaicing can be extended to capture not only a wide-angle image but also additional scene information. The key idea is illustrated on the left side of the figure, where we see a video camera with an optical filter with spatially varying properties attached to the front of the camera lens. In the example shown, the video camera is a black-and-white one and the filter is a linear interference one that passes a different wavelength of the visible light spectrum through each of its columns (see inset image). An image captured by the video camera is shown in the middle. The camera is moved with respect to a stationary scene and the acquired images are aligned using a registration algorithm. After registration, we have measurements of the radiance of each scene point for different wavelengths. These measurements are interpolated to obtain the spectral distribution of each scene point. The end result is the multispectral mosaic shown on the right side of Figure 2(c), instead of just a three-color (red, green, blue) mosaic that is obtained in the case of traditional mosaicing.

We refer to this approach as generalized mosaicing as it can be used to explore various dimensions of imaging by simply using the appropriate optical filter. A spatially varying neutral density filter may be used to capture a wide

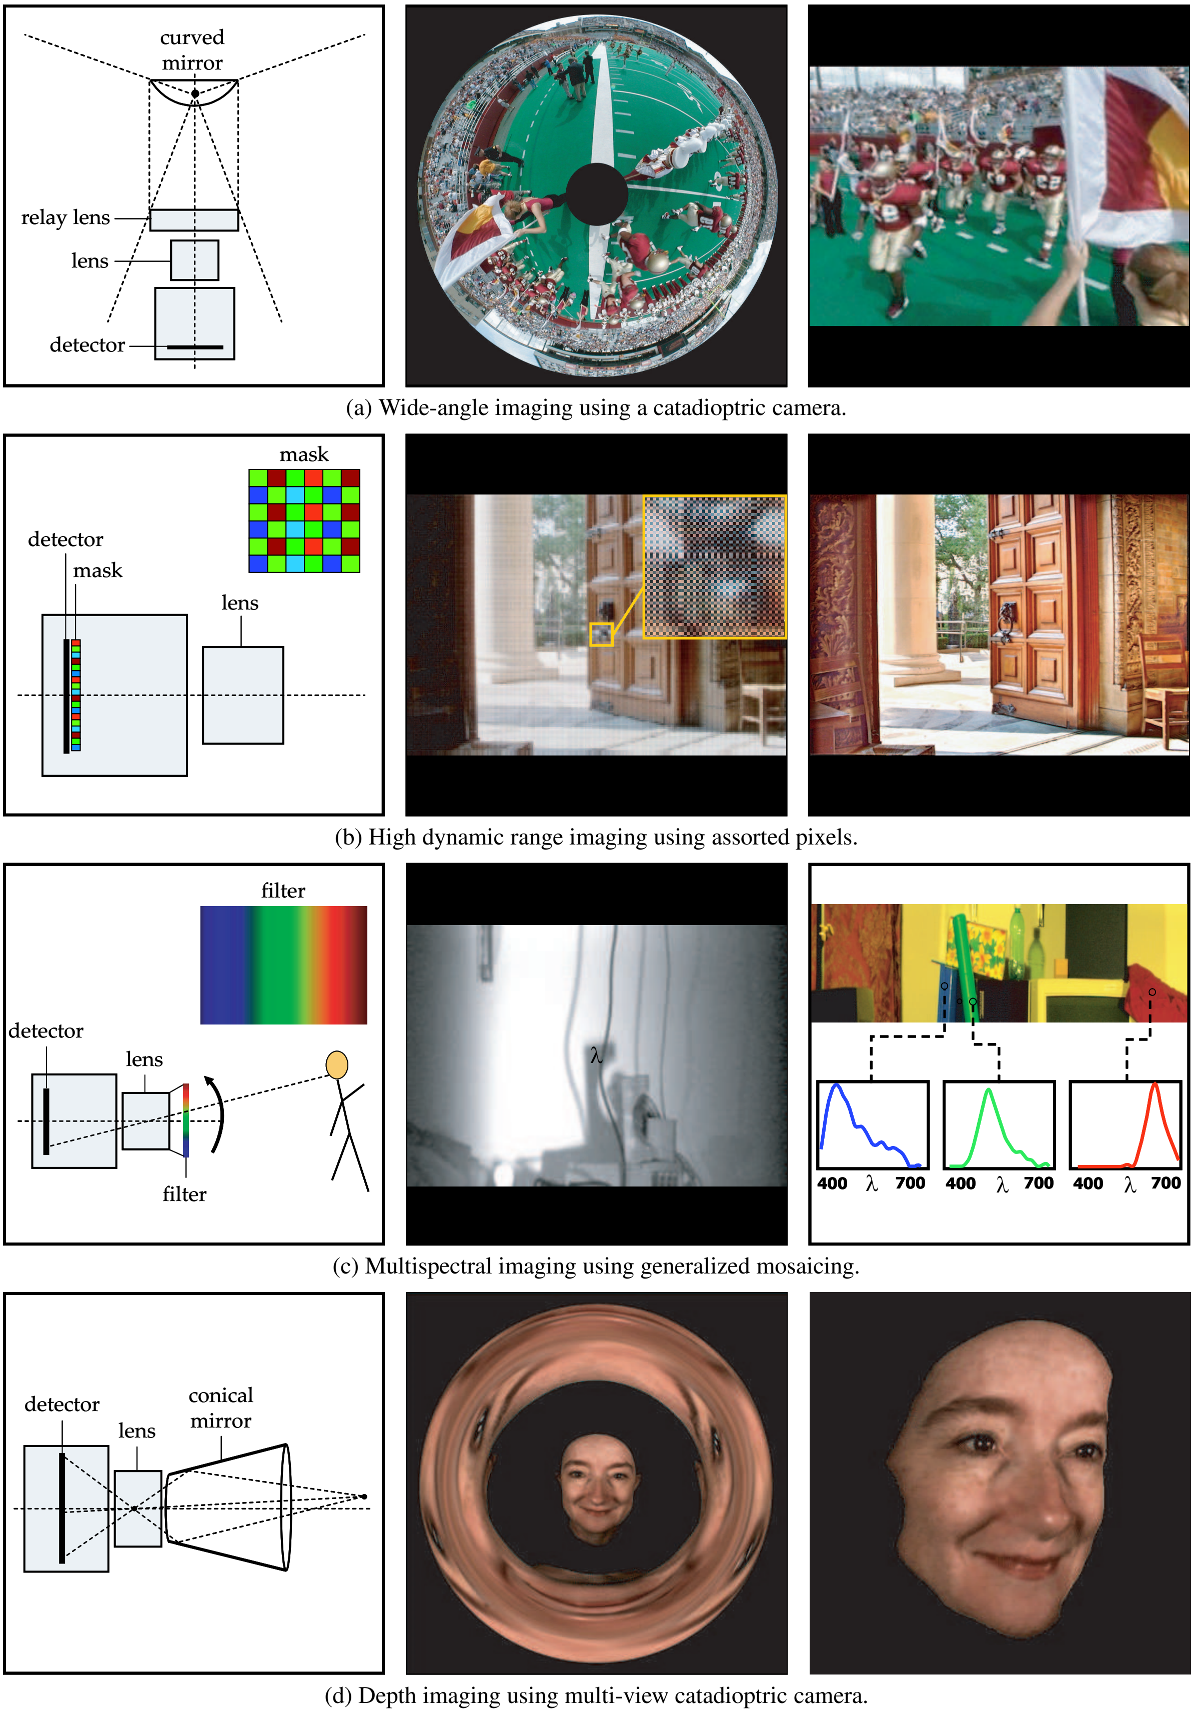

Figure 2. Examples of computational cameras that use unconventional optics and software to produce new types of images.

dynamic range mosaic and a filter with spatially varying polarization direction can be used to separate diffuse and specular reflections from the scene and detect material properties. When the filter is a wedge-shaped slab of glass, the scene points are measured under different focus settings and an all-focused mosaic can be computed. In fact, multiple imaging dimensions can be explored simultaneously by using more complex optical filters.

In Figure 2(d), we show how a computational camera can be used to extract the 3D structure of the scene from a single image. In front of a conventional perspective camera, we place a hollow cone that is mirrored on the inside. The axis of the cone is aligned with the optical axis of the camera. Since the mirror is hollow, a scene point is seen directly by the camera lens. In addition, it is reflected by exactly two points on the conical mirror that lie on a plane that passes through the scene point and the optical axis of the camera. As a result, each scene point is imaged from three different viewpoints: the center of projection of the camera lens and two virtual viewpoints that are equidistant and on opposite sides with respect to the optical axis. When one considers an entire scene, the image includes three views of it – one from the center of projection of the lens and two additional views from a circular locus of viewpoints whose center lies on the optical axis.

We refer to this type of a camera as a radial imaging system. An image of a face captured by the camera is shown in the middle of Figure 2(d). Notice how the center of the image is just a regular perspective view of the face. The annulus around this view has embedded within it two additional views of the face. A stereo matching algorithm is used to find correspondences between the three views and compute the 3D geometry of the face. The image on the right of Figure 2(d) shows a new rotated view of the face. While we used a conical mirror with specific parameters here, a variety of radial imaging systems with different imaging properties can be created by changing the parameters of the mirror. We have used this approach to recover the fine geometry of a 3D texture, capture complete texture maps of simple objects and measure the reflectance properties of real world materials.

3. Programmable Imaging

As we have seen, computational cameras produce images that are fundamentally different from the traditional perspective image. However, the hardware and software of each of these devices are designed to produce a particular type of image. The nature of this image cannot be altered without significant redesign of the device. This brings us to the notion of a programmable imaging system, which is illustrated in Figure 3. It uses an optical system for forming the image that can be varied by a controller in terms of its radiometric and/or geometric properties. When such a change is applied to the optics, the controller also changes the software in the computational module. The result is a single imaging system that can emulate the functionalities of several specialized ones. Such a flexible camera has two major benefits. First, a user is free to change the role of the camera based on his/her needs. Second, it allows us to explore the notion of a purposive camera that, as time progresses, always produces the visual information that is most pertinent to the task.

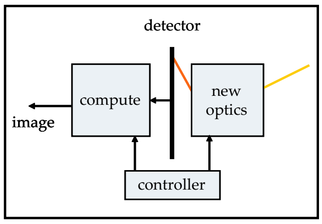

Figure 3. A programmable imaging system is a computational camera whose optics and software can be varied to emulate different imaging functionalities.

We now present two examples of programmable imaging systems. The first one, shown on the left of Figure 4(a), uses a two-dimensional array of micro-mirrors, whose orientations can be controlled. The image of the scene is first formed using a lens on the micro-mirror array. The plane on which the array resides is then re-imaged using a second lens onto an image detector. While it would be ideal to have a micro-mirror array whose mirror orientations can be set to any desired value, such a device is not available at this point in time. In our implementation, we have used the digital micro-mirror device (DMD) that has been developed by Texas Instruments [5] and serves as the workhorse for a large fraction of the digital projectors available today. The mirror of this array can only be switched between two orientations – 10 and -10 degrees. When a micro-mirror is oriented at 10 degrees the corresponding image detector pixel is exposed to a scene point and when it is at -10 degrees it receives no light. The switching between the two orientation states can be done in a matter of microseconds.

As an example, we show how this system can independently adapt the dynamic range of each of its pixels based on the brightness of the scene point it sees. In this case, the exposure of each pixel on the image detector is determined by the fraction of the integration time of the detector for which the corresponding micro-mirror on the DMD is oriented at 10 degrees. A simple control algorithm is used to update the exposure duration of each pixel based on the most recent captured image. The image in the middle of Figure 4(a) was captured by a conventional 8 bit video camera. The image on the right shows the output of the programmable imaging system with adaptive dynamic range. Note how the pixels that are saturated in the conventional camera image are brought into the dynamic range of the 8

(a) Adaptive dynamic range imaging with a micro-mirror array.

(b) Split field of view imaging with a volumetric aperture.

Figure 4. Programmable imaging systems that use controllable spatial light modulators to vary their radiometric and photometric properties based on the needs of the application.

bit camera. The inset image on the left of Figure 4(a) shows the adaptive exposure pattern applied to the micro-mirror array. This image can be used with the captured image on the right to compute an image with a very wide dynamic range. This imaging system has also been used to perform other imaging functionalities such as feature detection and object recognition.

In virtually any imaging system, the main reason to use a lens is to gather more light. As mentioned earlier, this benefit of a lens comes with the price that it severely restricts the geometric mapping of scene rays to image points. To address this limitation, we have been recently exploring lensless imaging systems. Consider a bare image detector exposed to a scene. In this case, each pixel on the detector receives a 2D set of rays of different directions from the scene. The detector itself is a 2D set of pixels of different spatial locations arranged on a plane. Therefore, although the detector produces a 2D image, it receives a 4D set of light rays from the scene. Now, consider a 3D (volumetric) aperture placed in front of the detector instead of a lens, as shown on the left of Figure 4(b). If the aperture has a 3D transmittance function embedded within it, it will modulate the 4D set of light rays before they are received by the 2D detector. If this transmittance function can be controlled, we would be able to apply a variety of modulation operations on the 4D set of rays. Such a device would enable us to map scene rays to pixels in ways that would be difficult, if not impossible, to achieve using a lens based camera.

Unfortunately, a controllable volumetric aperture is not easy to implement. Hence, we have implemented the aperture as a stack of controllable 2D apertures. Each aperture is a liquid crystal (LC) sheet of the type used in displays. By simply applying an image to the LC sheet, we can control its modulation function and change it from one captured image to the next. The inset image on the left of Figure 4(b) shows how three disconnected fields of view are projected onto adjacent regions on the detector, by appropriately selecting the open (full transmittance) and closed (zero transmittance) areas on two apertures. The advantage of such a “split field of view” projection is seen by comparing the middle and right images in Figure 4(b). The middle image was taken by a conventional camera. Although we are only interested in the three people in the scene, we are forced to waste a large fraction of the detector’s resolution on the scene regions in between the people. The right image was taken using the lensless system and we see that the three people are optically cropped out of the scene and imaged with higher resolution.

4. Programmable Illumination: A Smarter Flash

Since the dawn of photography people have been trying to take pictures of dimly lit scenes. The only way one could obtain a reasonably bright image of a dark scene was by using a very long exposure time, during which the scene had to remain stationary. The flashbulb was invented to overcome this limitation. The first commercial flashbulb appeared around 1930, and its design was based on patents awarded to a German inventor named Johannes Ostermeier. Today, the flashbulb, commonly referred to as the “flash,” is an integral part of virtually any consumer camera. In recent years, researchers have begun to explore ways to combine images taken with and without a flash to produce images of higher quality. Multiple flashes placed around the camera’s lens have also been used to detect depth discontinuities and produce stylized renderings of the scene.

It is interesting to note that the basic capability of the flash has remained the same since its invention. It is used to brightly illuminate the camera’s field of view during the exposure time of the image detector. It essentially serves as a point light source that illuminates everything within a reasonable distance from the camera. Given the enormous technological advancements made by digital projectors, the time may have arrived for the flash to play a more sophisticated role in the capture of images. The use of a projector-like source as a camera flash is powerful as it provides full control over the 2D set of rays it emits. It enables the camera to project arbitrarily complex illumination patterns onto the scene, capture the corresponding images, and compute information regarding the scene that is not possible to obtain with the traditional flash. In this case, the captured images are optically coded due to the patterned illumination of the scene.

We now present two examples that illustrate the benefits of using a digital projector as a programmable camera flash. On the left side of Figure 5(a), we see a camera and projector that are co-located by using a half-mirror. This configuration has the unique property that all the points that are visible to the camera can be illuminated by the projector. To maximize the brightness of the images they produce, projectors are made with large apertures and hence narrow depths of field. We have developed a method that exploits a projector’s narrow depth of field to recover the geometry of the scene viewed by the camera. The method uses a stripe pattern like the one shown in the inset image. This pattern is shifted a minimum of three times and the corresponding images are captured by the camera. The set of intensities measured at each camera pixel reveal the defocus of the shifted pattern, which in turn gives the depth of the scene point. This temporal defocus method has two advantages. First, since depth is computed independently for each camera pixel, it is able to recover sharp depth discontinuities. Second, since it is based on defocus and not triangulation, we are able to co-locate the projector and the camera and compute a depth map that is “image-complete,” i.e., there are no holes in the depth map from the perspective of the camera.

The middle of Figure 5(a) shows an image of a complex scene that includes a flower vase behind a wooden fence and its depth map (shown as gray-scale image) computed using the temporal defocus method. The depth map can be used to blur the scene image in a spatially varying manner to render an image as it would appear through a narrow depth of field camera lens. On the right of Figure 5(a) we see such a “refocused” image, where the petals in the back are in focus while the fence in the front is blurred. In short, a photographer can vary the depth of field of the image after it is captured. We have also used the depth maps computed using the temporal defocus method to insert synthetic objects within the captured image with all the desired occlusion effects.

Consider a scene lit by a point light source and viewed by a camera. The brightness of each scene point has two components, namely, direct and global. The direct component is due to light received by the point directly from the source and the global component is due to light received by the point from all other points in the scene. In our final example, we show how a programmable flash can be used to separate a scene into its direct and global components. The two components can then be used to edit the physical properties of objects in the scene and produce novel images.

Consider an image of the scene captured using the checkerboard illumination pattern shown in the inset image on the left of Figure 5(b). If the frequency of the checkerboard pattern is high, then the camera brightness of a point that is lit by one of the checkers includes the direct component and exactly half of the global component, since only half of the remaining scene points are lit by the checkerboard pattern. Now consider a second image captured using the complement of the above illumination pattern. In this case, the above scene point does not have a direct component but still produces exactly half of the global component. Since the above argument applies to all points in the scene, the direct and global components of all the scene points can be measured by projecting just two illumination patterns. In practice, to overcome the resolution limitations of the source, one may need to capture a larger set of images by shifting the checkerboard pattern in small steps.

In the middle of Figure 5(b), we show separation results for a scene with peppers of different colors. The direct image includes mainly the specular reflections from the surfaces of the peppers. The colors of the peppers come from subsurface scattering effects that are captured in the global image. This enables a user to alter the colors of the peppers

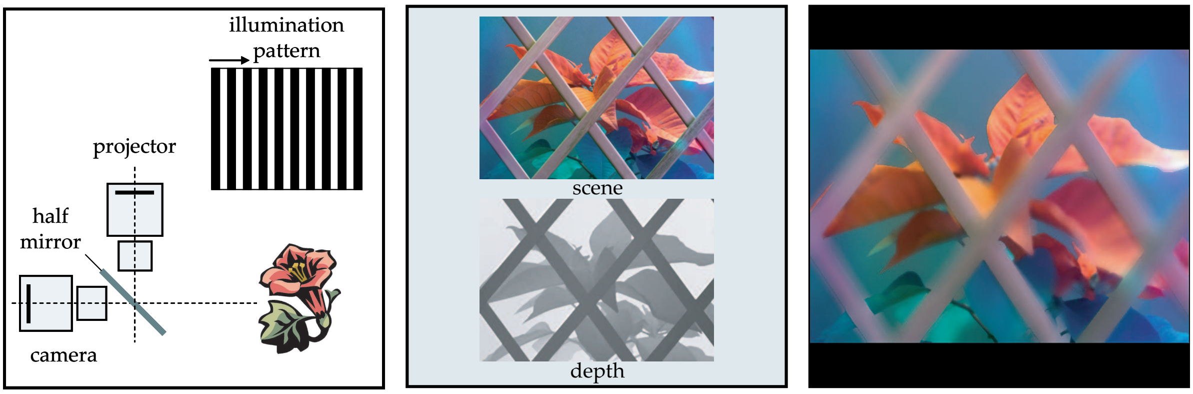

(a) Computing image-complete depth maps using projector defocus.

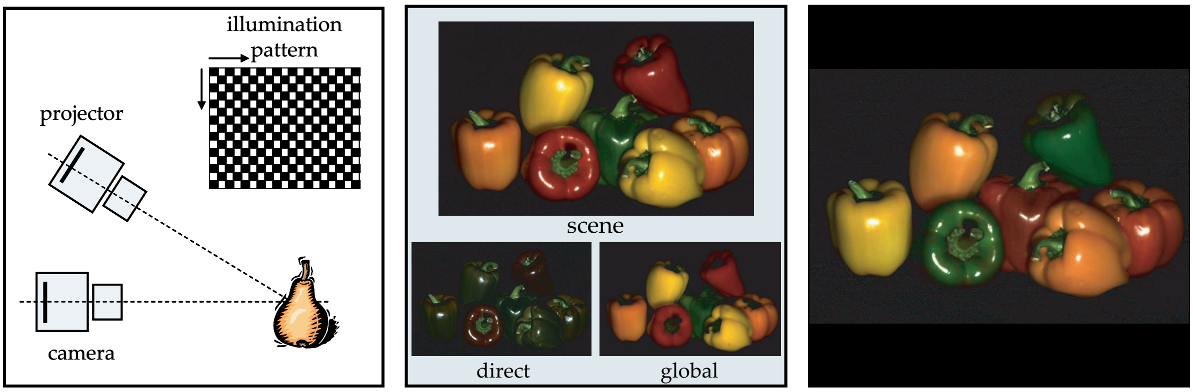

(b) Separation of direct and global illumination using high frequency illumination.

Figure 5. A projector can be used as a programmable camera flash to recover important scene information such as depth and illumination effects. Such information can be used to compute novel images of the scene.

in the global image and recombine it with the direct image to obtain a novel image, like the one shown on the right of Figure 5(b). In addition to subsurface scattering, the above separation method is applicable to other global illumination effects, including, interreflections between opaque surfaces and volumetric scattering from participating media.

5. Cameras of the Future

We have shown through examples how computational cameras use unconventional optics and software to produce new forms of visual information. We also described how this concept can be taken one step further by using controllable optics and software to realize programmable imaging systems that can change their functionalities based on the needs of the user or the application. Finally, we illustrated the benefits of using a programmable illumination source as a camera flash. Ultimately, the success of these concepts will depend on technological advances made in imaging optics, image detectors, and digital projectors. If progress in these fields continues at the remarkable pace we have seen in the last decade, we can expect the camera to evolve into a more versatile device that could further impact the ways in which we communicate with each other and express ourselves.

Acknowledgments

An earlier version of this article appeared in the August 2006 issue of IEEE Computer Magazine. The imaging systems described in this article were developed with support from NSF, ONR, DARPA, and the David and Lucile Packard Foundation. In particular, the author is grateful to Tom Strat, Tom McKenna and Behzad Kamgar-Parsi for their support and encouragement. The systems described here were developed by the author in collaboration with: Venkat Peri and Simon Baker (wide-angle imaging); Tomoo Mitsunaga (high dynamic range imaging); Yoav Schechner (generalized imaging); Sujit Kuthirummal (radial imaging); Vlad Branzoi and Terry Boult (programmable imaging with micro-mirror arrays); Assaf Zomet (programmable imaging with volumetric apertures); Li Zhang (projector defocus analysis); and Gurunandan Krishnan, Michael Grossberg and Ramesh Raskar (separation of direct and global illumination). The wide-angle image was provided by RemoteReality, Inc. The author thanks Anne Fleming, Gurunandan Krishnan and Li Zhang for their useful comments. Technical details on the imaging systems described in the article, and several others, can be found at *www.cs.columbia.edu/CAVE*.

References

[1] E. R. Dowski, Jr., and W. T. Cathey. Wavefront coding for detection and estimation with a single-lens incoherent optical system. In *Proceedings of ICASSP Conference*, volume 4, pages 2451–2454, May 1995.

[2] E. E. Fenimore and T. M. Cannon. Coded aperture imaging with uniformly redundant arrays. In *Applied Optics*, volume 17, pages 337–347, 1978.

[3] M. Foucault. *Discipline and Punish: The Birth of the Prison*. Vintage, New York, USA, 1979.

[4] A. Gershun. The Light Field. In *The Journal of Mathematics and Physics*, volume 18, pages 51–151. MIT, 1939.

[5] L. Hornbeck. Deformable-mirror spatial light modulators. In *Projection Displays III*, volume 1150, pages 86–102. SPIE, August 1989.

[6] T. Kanade and R. Bajcsy. Computational sensors. In *DARPA Workshop Report*, May 1993.

[7] G. Lippmann. La photographie intégral. In *Comptes-Rendus, Académie des Sciences*, volume 146, pages 446–551, 1908.

[8] P. L. Manly. *Unusual Telescopes*. Cambridge University Press, Cambridge, UK, 1991.

[9] C. Mead. *Analog VLSI and Neural Systems*. Addison-Wesley Longman Publishing Co., Boston, MA, USA, 1989.

[10] K. Miyamoto. Fish Eye Lens. *Journal of Optical Society of America*, 54(8):1060–1061, August 1964.

[11] S. K. Nayar. Computational cameras: Redefining the imaging. In *IEEE Computer Magazine*, pages 62–70, August 2006.

[12] B. Newhall. *The History of Photography*. The Museum of Modern Art, New York, 1982.

[13] J. L. Wyatt, C. Keast, M. Seidel, D. Standley, B. P. Horn, T. Knight, C. Sodini, H.-S. Lee, and T. Poggio. Analog VLSI Systems for Image Acquisition and Fast Early Vision Processing. In *International Journal of Computer Vision*, volume 8, pages 217–230, 1992.