

Computational Photography: Epsilon to Coded Photography

Ramesh Raskar

Media Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
raskar@media.mit.edu
<http://web.media.mit.edu/~raskar/>

Abstract. Computational photography combines plentiful computing, digital sensors, modern optics, actuators, and smart lights to escape the limitations of traditional cameras, enables novel imaging applications and simplifies many computer vision tasks. However, a majority of current Computational photography methods involves taking multiple sequential photos by changing scene parameters and fusing the photos to create a richer representation. Epsilon photography is concerned with synthesizing omnipictures and proceeds by multiple capture single image paradigm (MCSI). The goal of Coded computational photography is to modify the optics, illumination or sensors at the time of capture so that the scene properties are encoded in a single (or a few) photographs. We describe several applications of coding exposure, aperture, illumination and sensing and describe emerging techniques to recover scene parameters from coded photographs.

Keywords: Digital photography, Fourier transform, Fourier optics, Optical heterodyning, Coded aperture imaging, digital refocusing, plenoptic camera.

1 Introduction

Computational photography combines plentiful computing, digital sensors, modern optics, actuators, and smart lights to escape the limitations of traditional cameras, enables novel imaging applications and simplifies many computer vision tasks. 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 are just some of the new applications found in Computational photography.

In this paper, we discuss Coded photography which involves encoding of the photographic signal and post-capture decoding for improved scene analysis. With film-like photography, the captured image is a 2D projection of the scene. Due to limited capabilities of the camera, the recorded image is a partial representation of the view. Nevertheless, the captured image is ready for human consumption: what you see is what you almost get in the photo.

In Coded photography, the goal is to achieve a potentially richer representation of the scene during the encoding process. In some cases, Computational photography reduces to *Epsilon photography*, where the scene is recorded via multiple images, each captured by epsilon variation of the camera parameters. For example, successive images (or neighboring pixels) may have a different exposure, focus, aperture, view, illumination, or instant of capture. Each setting allows recording of partial information about the scene and the final image is reconstructed from these multiple observations. In Coded computational photography, the recorded image may appear distorted or random to a human observer. But the corresponding decoding recovers valuable information about the scene. *Less is more* in Coded photography. By blocking light over time or space, we can preserve more details about the scene in the recorded single photograph. In this paper we look at four specific examples:

1. **Coded exposure:** By blocking light in time, by fluttering the shutter open and closed in a carefully chosen binary sequence, we can preserve high spatial frequencies of fast moving objects to support high quality motion deblurring.
2. **Coded aperture optical heterodyning:** By blocking light near the sensor with a sinusoidal grating mask, we can record 4D light field on a 2D sensor. And by blocking light with a mask at the aperture, we can extend the depth of field and achieve full resolution digital refocussing.
3. **Coded illumination:** By observing blocked light at silhouettes, a multi-flash camera can locate depth discontinuities in challenging scenes without depth recovery.
4. **Coded sensing:** By sensing intensities with lateral inhibition, a gradient sensing camera can record large as well as subtle changes in intensity to recover a high-dynamic range image.

We describe several applications of Coding exposure, aperture, illumination and sensing and describe emerging techniques to recover scene parameters from coded photographs. But first, we give a introductory overview of the concepts involved in light fields.

1.1 What is a Light Field?

The light field is a function that describes the amount of light traveling in every direction through every point in space [17]. In geometric optics, the fundamental carrier of light is a ray. The measure for the amount of light traveling along a ray is radiance. The radiance along all such rays in a region of three-dimensional space illuminated by an unchanging arrangement of lights is called the *plenoptic* function. The plenoptic illumination function is an idealized function used in computer vision and computer graphics to express the image of a scene from any possible viewing position at any viewing angle at any point in time. Since rays in space can be parameterized by three spatial coordinates, x , y and z and two angles θ and ϕ it is a five-dimensional function [17].

The 4D Light Field. Radiance along a ray remains constant if there are no blockers. If we restrict ourselves to locations outside the convex hull of an




Fig. 1. A ray in 3D space is specified by its position (x, y, z) and direction (θ, ϕ)




Fig. 2. The two plane parametrization of the 4D light field: using pairs of points on two planes in any general position to represent the flow of light through an empty region of three-dimensional space [17]

object, then we can measure the plenoptic function easily using a digital camera. Moreover, in this case the function contains redundant information, because the radiance along a ray remains constant from point to point along its length. In fact, the redundant information is exactly one dimension, leaving us with a four-dimensional function. Parry Moon dubbed this function the photic field, while researchers in computer graphics call it the 4D light field or *Lumigraph* [12], [13]. Formally, the 4D light field is defined as radiance along rays in empty space.

Most commonly, the set of rays in a light field can be parameterized using the two-plane parametrization. While this parametrization cannot represent all rays, for example rays parallel to the two planes if the planes are parallel to each other, it has the advantage of relating closely to the analytic geometry of perspective imaging. A light field parameterized this way is sometimes called a light slab [17].

4D Reflectance Field. The bidirectional reflectance distribution function (BRDF) is a 4-dimensional function that defines how light is reflected at an opaque surface [18]. The function takes an incoming light direction, and outgoing direction, both defined with respect to the surface normal and returns the ratio of reflected radiance exiting along the outgoing direction to the irradiance incident on the surface from incoming direction. Note that each direction is itself




Fig. 3. When we measure all the light rays going out of the enclosure, it comprises of the 4D light field (figure from [16])

parameterized by azimuth angle and elevation, therefore the BRDF as a whole is 4-dimensional. As a further intuitive illustration [16] of 4D light fields imagine a convex enclosure of a 3D scene and an inward-facing ray camera at every surface point. Pick the *outgoing* rays you need for *any* camera outside the convex enclosure. The 2D surface of cameras and the 2D ray set for each camera gives rise to the 4D set of rays (4D light field of Lumigraph). When the similar idea is applied to the 4D set of *incoming* rays it comprises the 4D illumination field. Together, they give rise to the 8D reflectance field.

1.2 Film-Like Photography

Photography is the process of making pictures by, literally, *drawing with light* or recording the visually meaningful changes in the light leaving a scene. This goal was established for film photography about 150 years ago.

Currently, *digital photography* is electronically implemented film photography, refined and polished to achieve the goals of the classic film camera which were

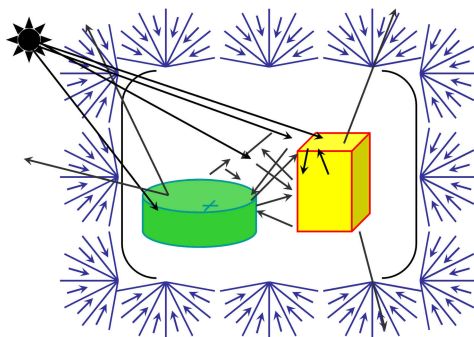


Fig. 4. When we measure all the light rays coming into the enclosure, it comprises of the 4D illumination field (figure from [16])

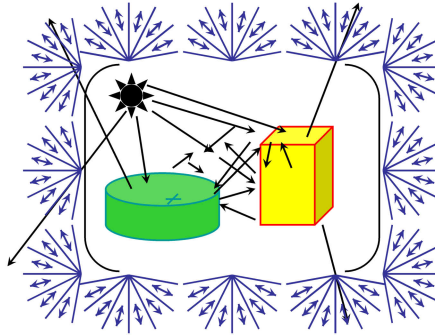


Fig. 5. Taken together, the 4D light field and the 4D illumination field give rise to the 8D reflectance field (Figure from [16]). Also define ratio $R_{ij} = \text{Outgoing ray}_i / \text{Incoming ray}_j$.

governed by chemistry, optics, mechanical shutters. Film-like photography presumes (and often requires) artful human judgment, intervention, and interpretation at every stage to choose viewpoint, framing, timing, lenses, film properties, lighting, developing, printing, display, search, index, and labelling.

In this article we plan to explore a progression away from film and film-like methods to something more comprehensive that exploits plentiful low-cost computing and memory with sensors, optics, probes, smart lighting and communication.

1.3 What Is Computational Photography?

Computational photography (CP) is an emerging field. We don't know where it will end up, we can't yet set its precise, complete definition, nor make a reliably comprehensive classification. But here is the scope of what researchers are currently exploring in this field.

- Computational photography attempts to record a richer visual experience, captures information beyond just a simple set of pixels and makes the recorded scene representation far more machine readable.
- It exploits computing, memory, interaction and communications to overcome long-standing limitations of photographic film and camera mechanics that have persisted in film-style digital photography, such as constraints on dynamic range, depth of field, field of view, resolution and the extent of scene motion during exposure.
- It enables new classes of recording the visual signal such as the *moment*, shape boundaries for non-photorealistic depiction [1], foreground versus background mattes, estimates of 3D structure, *relightable* photos and interactive displays that permit users to change lighting, viewpoint, focus, and more, capturing some useful, meaningful fraction of the *light field* of a scene, a 4-D set of viewing rays.

- It enables synthesis of impossible photos that could not have been captured at a single instant with a single camera, such as wrap-around views (*multiple-center-of-projection* images), fusion of time-lapsed events [1], the motion-microscope (motion magnification), video textures and panoramas. They also support impossible physical camera movements such as the *freeze effect* (in the movie Matrix) sequence recorded with multiple cameras with staggered exposure times.
- It encompass previously exotic forms of scientific imaging and data gathering techniques e.g. from astronomy, microscopy, and tomography.

1.4 Elements of Computational Photography

Traditional film-like photography involves a lens, a 2D planar sensor and a processor that converts sensed values into an image. In addition, the photography may involve external illumination from point sources (e.g. flash units) and area sources (e.g. studio lights). Computational photography generalizes the following four elements.

1. **Generalized optics:** Each optical element is treated as a 4D ray-bender that modifies a light field. The incident 4D light field for a given wavelength is transformed into a new 4D light field. The optics may involve more than one optical axis [15]. In some cases the perspective foreshortening of objects based on distance may be modified using wavefront coded optics [14]. In recent *lens-less* imaging methods and Coded aperture imaging used for gamma-ray and X-ray astronomy, the traditional lens is missing entirely. In some cases optical elements such as mirrors outside the camera adjust the linear combinations of ray bundles that reach the sensor pixel to adapt the sensor to the viewed scene.
2. **Generalized sensors:** All light sensors measure some combined fraction of the 4D light field impinging on it, but traditional sensors capture only a 2D projection of this light field. Computational photography attempts to capture more; a 3D or 4D ray representation using planar, non-planar or even volumetric sensor assemblies. For example, a traditional out-of-focus 2D image is the result of a capture-time decision: each detector pixel gathers light from its own bundle of rays that do not converge on the focused object. But a plenoptic Camera [9], [10] subdivides these bundles into separate measurements. Computing a weighted sum of rays that converge on the objects in the scene creates a digitally refocused image, and even permits multiple focusing distances within a single computed image. Generalizing sensors can extend their dynamic range [2] and wavelength selectivity as well. While traditional sensors trade spatial resolution for color measurement (wavelengths) using a Bayer grid or red, green or blue filters on individual pixels, some modern sensor designs determine photon wavelength by sensor penetration, permitting several spectral estimates at a single pixel location.

3. **Generalized reconstruction:** Conversion of raw sensor outputs into picture values can be much more sophisticated. While existing digital cameras perform *de-mosaicking*, (interpolate the Bayer grid), remove fixed-pattern noise, and hide *dead* pixel sensors, recent work in computational photography can do more. Reconstruction might combine disparate measurements in novel ways by considering the camera intrinsic parameters used during capture. For example, the processing might construct a high dynamic range scene from multiple photographs from coaxial lenses, from sensed gradients, [2] or compute sharp images of a fast moving object from a single image taken by a camera with a *fluttering* shutter [3]. Closed-loop control during photography itself can also be extended, exploiting traditional cameras' exposure control, image stabilizing, and focus, as new opportunities for modulating the scene's optical signal for later decoding.
4. **Computational illumination:** Photographic lighting has changed very little since the 1950's: with digital video projectors, servos, and device-to-device communication, we have new opportunities to control the sources of light with as much sophistication as we use to control our digital sensors. What sorts of spatio-temporal modulations for light might better reveal the visually important contents of a scene? Harold Edgerton showed high-speed strobes offered tremendous new appearance-capturing capabilities; how many new advantages can we realize by replacing the *dumb* flash units, static spot lights and reflectors with actively controlled spatio-temporal modulators and optics? Already we can capture occluding edges with multiple flashes [1], exchange cameras and projectors by Helmholtz reciprocity, gather relichtable actor's performances with light stages and see through muddy water with coded-mask illumination. In every case, better lighting control during capture allows one to build richer representations of photographed scenes.

2 Sampling Dimensions of Imaging

2.1 Epsilon Photography for Optimizing Film-Like Cameras

Think of film cameras at their best as defining a *box* in the multi-dimensional space of imaging parameters. The first, most obvious thing we can do to improve digital cameras is to expand this box in every conceivable dimension. This effort reduces Computational photography to *Epsilon photography*, where the scene is recorded via multiple images, each captured by epsilon variation of the camera parameters. For example, successive images (or neighboring pixels) may have different settings for parameters such as exposure, focus, aperture, view, illumination, or the instant of capture. Each setting allows recording of partial information about the scene and the final image is reconstructed from these multiple observations. Epsilon photography is thus concatenation of many such boxes in parameter space; multiple film-style photos computationally merged to make a more complete photo or scene description. While the merged photo is superior, each of the individual photos is still useful and comprehensible on its

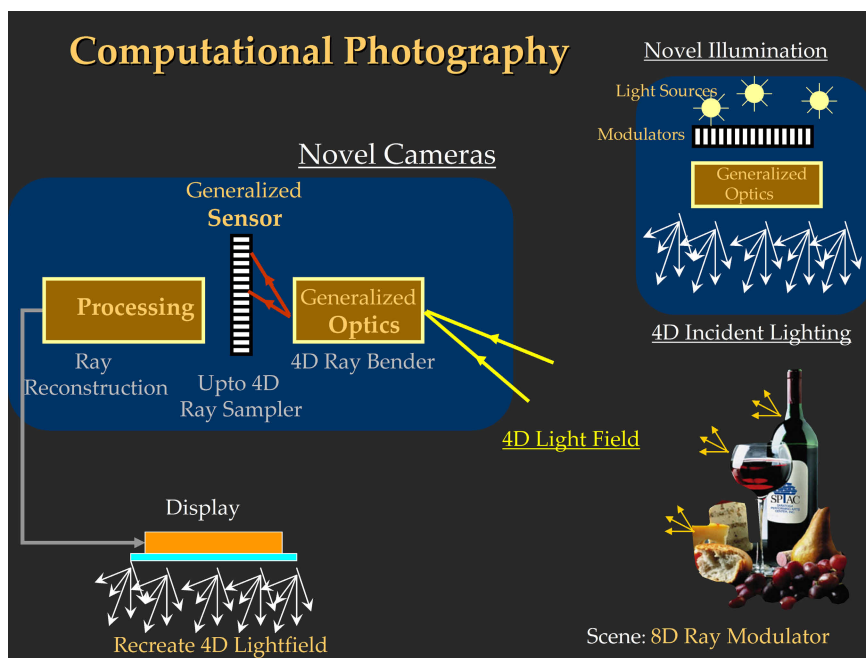


Fig. 6. Elements of Computational photography

own, without any of the others. The merged photo contains the best features from all of them.

1. **Field of view:** A wide field of view panorama is achieved by stitching and mosaicking pictures taken by panning a camera around a common center of projection or by translating a camera over a near-planar scene.
2. **Dynamic range:** A high dynamic range image is captured by merging photos at a series of exposure values [6]
3. **Depth of field:** All-in-focus image is reconstructed from images taken by successively changing the plane of focus.
4. **Spatial resolution:** Higher resolution is achieved by tiling multiple cameras (and mosaicking individual images) or by jittering a single camera.
5. **Wavelength resolution:** Traditional cameras sample only 3 basis colors. But multi-spectral (multiple colors in the visible spectrum) or hyper-spectral (wavelengths beyond the visible spectrum) imaging is accomplished by taking pictures while successively changing color filters in front of the camera, using tunable wavelength filters or using diffraction gratings.
6. **Temporal resolution:** High speed imaging is achieved by staggering the exposure time of multiple low-frame rate cameras. The exposure durations of individual cameras can be non-overlapping or overlapping.

Taking multiple images under varying camera parameters can be achieved in several ways. The images can be taken with a single camera over time. The

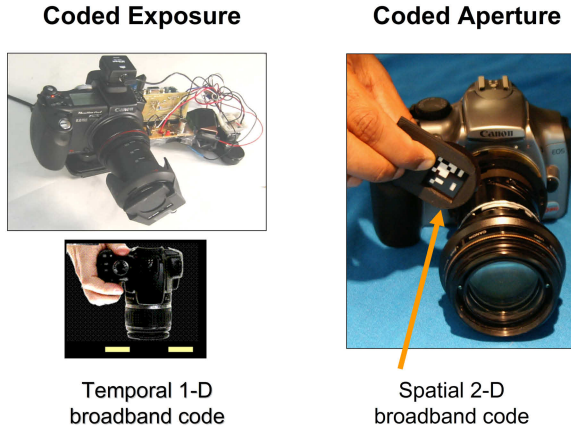


Fig. 7. Blocking light to achieve Coded photography. (Left) Using a 1-D code in time to block and unblock light over time, a coded exposure photo can reversibly encode motion blur ([3]). (Right) Using a 2-D code in space to block parts of the light via a masked aperture, a coded aperture photo can reversibly encode defocus blur ([4]).

images can be captured simultaneously using *assorted pixels* where each pixel is a tuned to a different value for a given parameter [5]. Simultaneous capture of multiple samples can also be recorded using multiple cameras, each camera having different values for a given parameter. Two designs are currently being used for multi-camera solutions: a camera array and single-axis multiple parameter (co-axial) cameras [8].

2.2 Coded Photography

There is much more beyond the *best possible film camera*. We can virtualize the notion of the camera itself if we consider it as a device that collects bundles of rays, each ray with its own wavelength spectrum and exposure duration.

Coded photography is a notion of an *out-of-the-box* photographic method, in which individual (ray) samples or data sets may or may not be comprehensible as *images* without further decoding, re-binning or reconstruction. Coded aperture techniques, inspired by work in astronomical imaging, try to preserve high spatial frequencies so that out of focus blurred images can be digitally re-focused [4]. By coding illumination, it is possible to decompose radiance in a scene into direct and global components. Using a Coded exposure technique, one can rapidly flutter open and close the shutter of a camera in a carefully chosen binary sequence, to capture a single photo. The fluttered shutter encoded the motion in the scene in the observed blur in a reversible way. Other examples include confocal images and techniques to recover glare in the images.

We may be converging on a new, much more capable *box* of parameters in computational photography that we don't yet recognize; there is still quite a bit







Coding in Time Exposure	Coding in Space Mask in aperture	Coded Illumination Inter-view	Coded Sensing Gradient Sensor
			
[Raskar et al 2006]	[Veeraraghavan et al 07]	[Raskar et al 2004]	[Tumblin et al 2005]
Optical Heterodyning Mask near sensor		Coded Illumination Intra-view	
			
[Veeraraghavan et al 07]		[Nayar et al 2006]	

Fig. 8. An overview of projects. Coding in time or space, coding the incident active illumination and coding the sensing pattern.

of innovation to come! In the rest of the article, we survey recent techniques that exploit exposure, focus, active illumination and sensors.

3 Coded Exposure

In a conventional single-exposure photograph, moving objects or moving cameras cause motion blur. The exposure time defines a temporal box filter that smears the moving object across the image by convolution. This box filter destroys important high-frequency spatial details so that deblurring via deconvolution becomes an ill-posed problem. We have proposed to flutter the camera's shutter open and closed during the chosen exposure time with a binary pseudo-random sequence, instead of leaving it open as in a traditional camera [3]. The flutter changes the box filter to a broad-band filter that preserves high-frequency spatial details in the blurred image and the corresponding deconvolution becomes a well-posed problem.

Results on several challenging cases of motion-blur removal including outdoor scenes, extremely large motions, textured backgrounds and partial occluders were presented. However, the authors assume that PSF (Point spread function) is given or is obtained by simple user interaction. Since changing the integration time of conventional CCD cameras is not feasible, an external ferro-electric shutter is placed in front of the lens to code the exposure. The shutter is driven opaque and transparent according to the binary signals generated from PIC [20] using the pseudo-random binary sequence.




Fig. 9. The flutter shutter camera. The Coded exposure is achieved by fluttering the shutter open and closed. Instead of a mechanical movement of the shutter, we used a ferro-electric LCD in front of the lens. It is driven opaque and transparent according to the desired binary sequence.

4 Coded Aperture and Optical Heterodyning

Can we capture additional information about a scene by inserting a patterned mask inside a conventional camera? We use a patterned attenuating mask to encode the light field entering the camera. Depending on where we put the mask, we can effect desired frequency domain modulation of the light field. If we put the mask near the lens aperture, we can achieve full resolution digital refocussing. If we put the mask near the sensor, we can recover a 4D light field without any additional lenslet array.

Coded aperture imaging has been historically used in radar (SAR) [19]. Ren Ng et. al. have developed a camera that can capture the 4D light field incident on the image sensor in a single photographic exposure [10]. This is achieved by inserting a microlens array between the sensor and main lens, creating a plenoptic camera. Each microlens measures not just the total amount of light deposited at that location, but how much light arrives along each ray. By re-sorting the measured rays of light to where they would have terminated in slightly different, synthetic cameras, one can compute sharp photographs focused at different depths. A linear increase in the resolution of images under each microlens results in a linear increase in the sharpness of the refocused photographs. This property allows one to extend the depth of field of the camera without reducing the aperture, enabling shorter exposures and lower image noise.




Fig. 10. Encoded blur camera, i.e. with mask in the aperture, can preserve high spatial images frequencies in the defocus blur. Notice the glint in the eye. In the misfocused photo, on the left, the bright spot appears blurred with the bokeh [21] of the chosen aperture (shown in the inset). In the deblurred result, on the right, the details on the eye are correctly recovered.

Our group has shown that it is also possible to create a plenoptic camera using a patterned mask instead of a lenslet array. The geometric configurations remains nearly identical [4]. The method is known as *spatial optical heterodyning*. Instead of remapping rays in 4D using microlens array so that they can be captured on a 2D sensor, spatial optical heterodyning remaps frequency components of the 4D light field so that the frequency components can be recovered from Fourier transform of the captured 2D image. In microlens array based design, each pixel effectively records light along a single ray bundle. With patterned masks, each pixel records a linear combination multiple ray-bundles. By carefully coding the linear combination, the coded heterodyning method can reconstruct the values of individual ray-bundles.

This is reversible modulation of 4D light field by inserting a patterned planar mask in the optical path of a lens based camera. We can reconstruct the 4D light field from a 2D camera image. The patterned mask attenuates light rays inside the camera instead of bending them, and the attenuation recoverably encodes the ray on the 2D sensor. Our mask-equipped camera focuses just as a traditional camera might to capture conventional 2D photos at full sensor resolution, but the raw pixel values also hold a modulated 4D light field. The light field can be recovered by rearranging the tiles of the 2D Fourier transform of sensor values into 4D planes, and computing the inverse Fourier transform.

5 Coded Illumination

By observing blocked light at silhouettes, a multi-flash camera can locate depth discontinuities in challenging scenes without depth recovery. We used a multi-flash camera to find the silhouettes in a scene [1]. We take four photos of an object with four different light positions (above, below, left and right of the

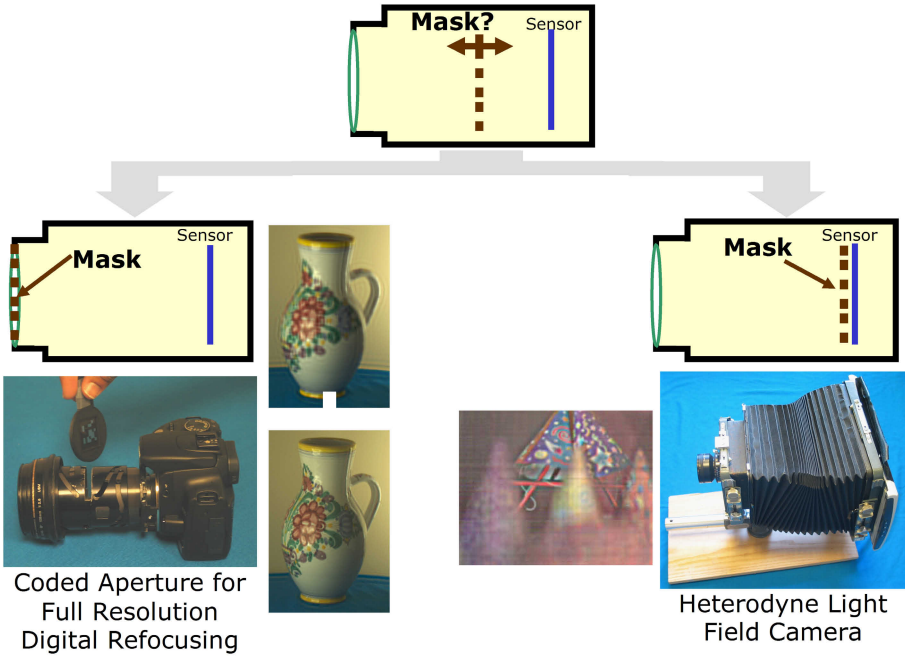


Fig. 11. Coding Light Field entering a camera via a mask

lens). We detect shadows cast along the depth discontinuities and use them to detect depth discontinuities in the scene. The detected silhouettes are then used for stylizing the photograph and highlighting important features. We also demonstrate silhouette detection in a video using a repeated fast sequence of flashes.

6 High Dynamic Range Using a Gradient Camera

A camera sensor is limited in the range of highest and lowest intensities it can measure. To capture the high dynamic range, one can adaptively set the exposure the sensor so that the signal to noise ratio is high over the entire image, including in the the dark and brightly lit regions. One approach for faithfully recording the intensities in a high dynamic range scenes is to capture multiple images using different exposures, and then to merge these images. The basic idea is that when longer exposures are used, dark regions are well exposed but bright regions are saturated. On the other hand, when short exposures are used, dark regions are too dark but bright regions are well imaged. If exposure varies and multiple pictures are taken of the same scene, value of a pixel can be taken from those images where it's neither too dark nor saturated. This type of approach is often referred to as exposure bracketing.

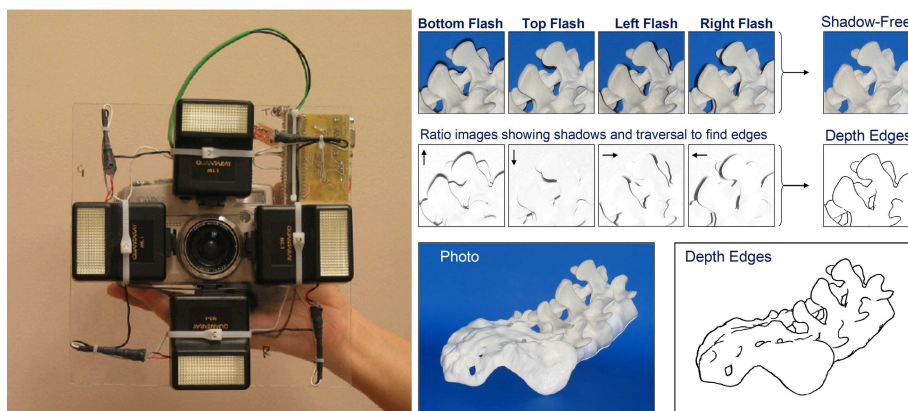


Fig. 12. Multi-flash Camera for Depth Edge Detection. (Left) A camera with four flashes. (Right) Photos due to individual flashes, highlighted shadows and epipolar traversal to compute the single pixel depth edges.

At the sensor level, various approaches have also been proposed for high dynamic range imaging. One type of approach is to use multiple sensing elements with different sensitivities within each cell. Multiple measurements are made from the sensing elements, and they are combined on-chip before a high dynamic range image is read out from the chip. Spatial sampling rate is lowered in these sensing devices, and spatial resolution is sacrificed. Another type of approach is to adjust the well capacity of the sensing elements during photocurrent integration but this gives higher noise.

By sensing intensities with lateral inhibition, a gradient sensing camera can record large as well as subtle changes in intensity to recover a high-dynamic range image. By sensing difference between neighboring pixels instead of actual intensities, our group has shown that a *Gradient Camera* can record large global variations in intensity [2]. Rather than measure absolute intensity values at each pixel, this proposed sensor measures only forward differences between them, which remain small even for extremely high-dynamic range scenes, and reconstructs the sensed image from these differences using Poisson solver methods. This approach offers several advantages: the sensor is nearly impossible to over- or under-expose, yet offers extremely fine quantization, even with very modest A/D convertors (e.g. 8 bits). The thermal and quantization noise occurs in the gradient domain, and appears as low frequency *cloudy* noise in the reconstruction, rather than uncorrelated high-frequency noise that might obscure the exact position of scene edges.

7 Conclusion

As these examples indicate, we have scarcely begun to explore the possibilities offered by combining computation, 4D modeling of light transport, and

novel optical systems. Nor have such explorations been limited to photography and computer graphics or computer vision. Microscopy, tomography, astronomy and other optically driven fields already contain some ready-to-use solutions to borrow and extend. If the goal of photography is to capture, reproduce, and manipulate a meaningful visual experience, then the camera is not sufficient to capture even the most rudimentary birthday party. The human experience and our personal viewpoint is missing. Computational Photography can supply us with visual experiences, but can't decide which one's matter most to humans. Beyond coding the first order parameters like exposure, focus, illumination and sensing, maybe the ultimate goal of Computational Photography is to encode the human experience in the captured single photo.

Acknowledgements

We wish to thank Jack Tumblin and Amit Agrawal for contributing several ideas for this paper. We also thank co-authors and collaborators Ashok Veeraraghavan, Ankit Mohan, Yuanzen Li, Karhan Tan, Rogerio Feris, Jingyi Yu, Matthew Turk. We thank Shree Nayar and Marc Levoy for useful comments and discussions and Ahmed Kirmani for a thorough rewrite and editing of this final draft.

References

1. Raskar, R., Tan, K., Feris, R., Yu, J., Turk, M.: Non-photorealistic Camera: Depth Edge Detection and Stylized Rendering Using a Multi-Flash Camera. In: Proc. ACM SIGGRAPH (2004)
2. Tumblin, J., Agrawal, A., Raskar, R.: Why I want a Gradient Camera. IEEE Computer Vision and Pattern Recognition (2005)
3. Raskar, R., Agrawal, A., Tumblin, J.: Coded exposure photography: motion deblurring using fluttered shutter. ACM Transactions on Graphics 25(3), 795–804 (2006)
4. Veeraraghavan, A., Raskar, R., Agrawal, A., Mohan, A., Tumblin, J.: Dappled Photography: Mask-Enhanced Cameras for Heterodyned Light Fields and Coded Aperture Refocusing. In: Proc. ACM SIGGRAPH (2007)
5. Nayar, S.K., Narasimhan, S.G.: Assorted Pixels: Multi-Sampled Imaging With Structural Models. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2353, pp. 636–652. Springer, Heidelberg (2002)
6. Debevec, M.: Recovering high dynamic range radiance maps from photographs. In: Proc. ACM SIGGRAPH (1997)
7. Mann, P.: Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In: Proc. Imaging Science and Technology 46th ann. conference (1995)
8. Morgan, M., Matusik, P., Hughes, D.: Defocus Video Matting. ACM Transactions on Graphics 24(3) (July 2005) (Proceedings of ACM SIGGRAPH 2005)
9. Adelson, E.H., Wang, J.Y.A.: Single Lens Stereo with a Plenoptic Camera. IEEE Transactions on Pattern Analysis and Machine Intelligence 14(2) (February 1992)
10. Ren, N.: Fourier Slice Photography. In: ACM SIGGRAPH (2005)

11. Morimura: Imaging method for a wide dynamic range and an imaging device for a wide dynamic range. U.S. Patent 5455621 (October 1993)
12. Levoy, M., Hanrahan, P.: Light field rendering. In: ACM SIGGRAPH, pp. 31–42 (1996)
13. Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F.: The Lumigraph. In: ACM SIGGRAPH, pp. 43–54 (1996)
14. Dowski Jr., E.R., Cathey, W.T.: Extended depth of field through wave-front coding. *Applied Optics* 34(11), 1859–1866 (1995)
15. Georgiev, T., Zheng, C., Nayar, S., Salesin, D., Curless, B., Intwala, C.: Spatio-angular Resolution Trade-Offs in Integral Photography. In: Proceedings of Eurographics Symposium on Rendering (2006)
16. Tumblin, J.: Slides on the Photographic Signal and Film-like Photography. In: Course 3: Computational Photography, ACM SIGGRAPH (2005), www.merl.com/people/raskar/photo/Slides/01BasicJTJuly31.ppt
17. Light fields, http://en.wikipedia.org/wiki/Light_field
18. Bidirectional reflectance distribution function, <http://en.wikipedia.org/wiki/BRDF>
19. Synthetic Aperture Radar, http://en.wikipedia.org/wiki/Synthetic_aperture_radar
20. Programmable Interface Controller, http://en.wikipedia.org/wiki/PIC_microcontroller
21. Bokeh, <http://en.wikipedia.org/wiki/Bokeh>