Title: The Future of Photography is Computational Photography

Subtitle: 100 years after invention, Integral Photography becomes feasible

Adobe customers are creative and demanding. They expect to use our Adobe Photoshop software to create the "impossible photograph." They want what they *imagined*, not the *image* that the camera captured. To help them close the gap between the image and their imaginations, we are constantly researching and designing new approaches.

The computational power available to Adobe Photoshop has grown exponentially, thanks to Moore's Law. We at Adobe find that we now have algorithms for image manipulation and enhancement that cannot be applied commercially without advances in the camera as well. This fact has led us to innovate not only in software, but also in hardware, and we've begun prototyping the kind of computational cameras and lenses we would like to see produced in order to enable these new algorithms. The integral (or plenoptic) camera is a promising advancement in this work.

In this article, we will outline our motivations and inspirations for this research. We will show how our work builds on that of other scientists, especially Gabriel Lippmann, winner of the Nobel Prize for color photography. We will discuss his idea of the "integral photograph" (Lippmann 1908) and how current research manifests this idea. We also will describe the applications of this novel approach, and demonstrate some of the results of our research on the merging of photography and computation. Perhaps we may even inspire you to contribute in some way to the future of photography, which we believe is computational photography.

Integral photography, past and future

The basic technology of capturing an image as a photograph has not changed since the inception of the camera; the visual field of light perceived three-dimensionally is rendered onto a two-dimensional plane – a photographic plate, film, or pixel grid. Meanwhile, computational power has increased exponentially. Because of new techniques in computer imaging software like Photoshop, photographers can achieve images that were previously either unattainable with a camera and darkroom, or prohibitively labor intensive. But we can achieve more; and our customers expect more. Today we are drawing inspiration for what's possible in the future of photography from great ideas of the past.

In 1908, Lippmann introduced a method of photographically capturing radiance – the physical quantity that describes the flow of light energy that exists in three-dimensional space. He called the recorded data "integral photographs." If we can make Lippmann's integral photography practical for mainstream photographers, it will be possible to do entirely new things in the "digital darkroom" that is Adobe Photoshop. With new algorithms and improvements in camera hardware like the integral camera, photographers can refocus the image, manipulate camera viewpoint, and employ other three-dimensional features like controlling focus on a per-pixel basis – all **after** the image has been captured. At Adobe we constantly look to the future and ask, "Why not?" For instance, Photoshop already has a Healing Brush which allows people to "heal" blemishes in a photograph. Why not a "Focus Brush" that allows users to paint any part of an image into or out of focus? Why not a new kind of Selection Brush that "knows" the depth of every pixel in an image, and thus enables users to quickly select and manipulate just parts of an image based on depth? These and other manipulations are not practical with today's cameras. Integral cameras and Photoshop could allow photographers to create with ease what has previously been an impossible photograph.

How does it work? An explanation involving "deep pixels"

Imagine a hypothetical camera with layers of pixels instead of a single plane of pixels. Light rays passing through the lens of a camera hit the sensor from different angles. The physical quantity describing the amount of energy flowing along the rays is the four-dimensional **radiance**, a function of the position and angle of the rays as they hit the sensor. Typically, this radiance information is flattened onto pixels to create a final two-dimensional image that is a function of the depth of field and focus of an image element at the moment of exposure. A “deep pixel,” however, would record information about the light rays at various depths. As light passes through the pixel layers, each layer would capture the focused image only from a given depth in the scene. These layers could then be recombined selectively in order to produce the desired focus; we could either select a given layer to be our final image or combine in-focus layers to create an all-in-focus image.

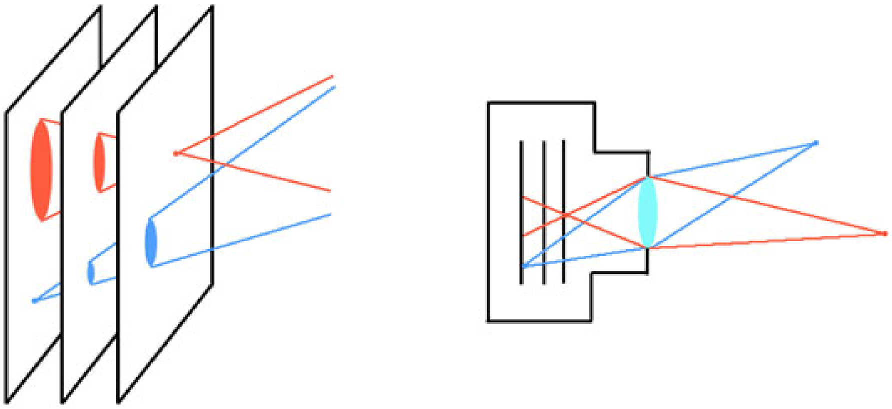

Figure 1: Our imaginary deep pixel camera.

The concept of the deep pixel is achieved optically in integral photography. Instead of a single lens focusing an image onto the image plane of a camera, integral photography employs multiple lenses which focus multiple images onto the image plane. In this way, the rays of light that pass through three-dimensional space are captured individually as angular information in the camera’s sensors.

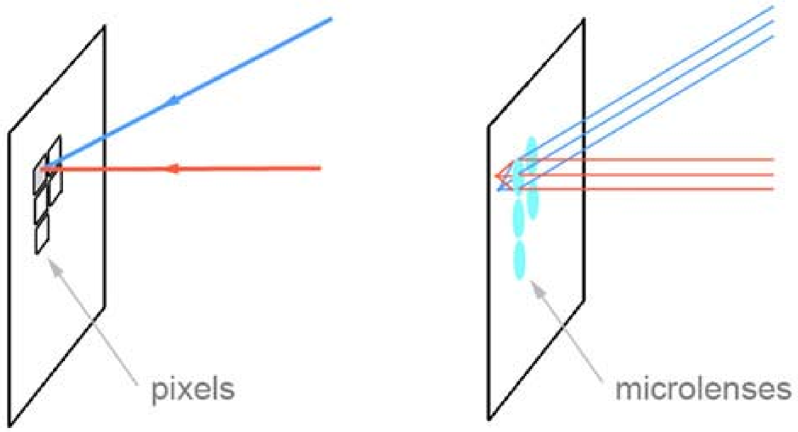

Figure 2. In a typical digital camera (left), radiance information is flattened on a pixel. An integral or plenoptic camera (right) captures the radiance via the focusing action of micro-lenses.

With this approach, a conventional camera pixel would be divided into multiple sub-pixels which individually capture rays from distinct directions. Thus a four-dimensional image is formed, a record of the full four-dimensional radiance at the sensor. Rays of light passing through a plane in three-dimensional space are stored as four-dimensional data (x,y, α, β). The position on the plane through which a ray passes (x,y) and the direction of the ray (α, β) make up the optical phase space (ray space), as in Figure 2.

After computer processing, the pixels act as if they are “deep,” rendering an image with virtual depth that the photographer can work with after the moment of image capture has passed. Our image processing with this technology effectively replaces the camera’s focusing lens with a virtual lens. Focus, lens aberration correction, and optics in general become functions of the software used to render the radiance information captured by the camera. This would enable entirely new ways of thinking about what a “lens” is: if we can correct for lens aberrations in the digital darkroom, instead of having to manufacture perfect lenses, we can have high quality images with inexpensive and even non-standard lenses.

Tradeoffs between angular and spatial resolution – can we have both?

Traditional approaches to capturing four-dimensional data in two dimensions require multiplexing by allocation of both angular and spatial information across a given two-dimensional pixel array. Thus there are tradeoffs: the result is either more angular information (greater depth of field) and fewer spatial pixels (lower resolution), or vice versa. (Insert footnote 1 a and b) At Adobe, we are researching better ways of optically multiplexing the four dimensional radiance onto the two-dimensional image plane to achieve precise control over how depth of field and resolution are represented in the final image. Here is an overview of two approaches and our improvements. The four-dimensional ray space can be split into spatial and angular components in two ways. One way is by placing an array of lenses (or lenses and prisms) in front of the camera. Using this method, researchers at Adobe Labs partitioned the optical phase space into 19 pictures capturing the scene from different vantage points (Georgiev et al. 2006). Another way to do this is by placing micro-lenses directly in front of the image sensor as proposed by

Deleted: ¶

Lippmann himself (Lippmann 1908), so that the image behind each micro-lens acts like the deep pixel we described earlier.

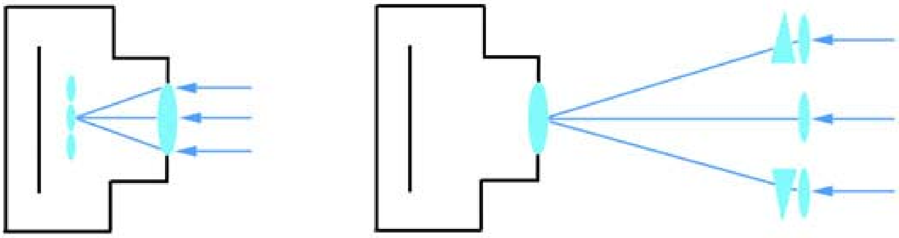

Figure 3. Two optical recipes for capturing the radiance. On the left, micro-lenses are placed one focal length in front of the sensor. The main camera lens forms an image at the plane of the micro-lenses. On the right, external lenses create an array of images in front of the main lens. The main lens is focused on this array of images, which is then reimaged to the sensor.

One problem with both designs is that they are fixed; the number of (micro-) lenses defines the number of pieces into which we can partition the ray space. Given that the number of pixels available to be divided is finite, lens developers have to make concessions of either depth or resolution as they attempt to define the most efficient partition of pixels in the spatio-angular tradeoff.

This tradeoff is related to another issue we encounter in integral photography. The physical world is primarily made of objects with Lambertian surfaces; all rays coming from one point on the surface of the object have essentially the same radiance. Because of this property, radiance data tends to have a redundancy in the angular dimension.

To resolve these complementary issues requires a new approach to focusing the plenoptic camera that takes angular redundancy into account. Our solution is a flexible partition of the optical phase space with a movable tradeoff point between spatial and angular resolution. Adjusting this tradeoff point to make use of the above angular redundancy, we can capture very low angular resolution radiance information to produce a high spatial resolution image.

Deleted: requires a two-fold approach

Deleted: We can make

Deleted: ; and

Deleted: , making use of redundancy to

Can we make integral photography practical?

The availability of ever more powerful computers puts us in a unique position in the hundred-year history of integral photography. Adelson and Wang of MIT have proposed the Plenoptic camera, a digital implementation of Lippmann's camera with application to computer vision (Adelson and Wang 1992). Further research has introduced algorithms for refocusing and other three-dimensional manipulation of radiance captured by the plenoptic or light field camera (Levoy and Hanrahan 1996, Ng et al. 2005). But even with these achievements, camera hardware is still a limiting factor. Is there a way to modify the plenoptic camera to leverage the computational power we have today and to make integral photography practical?

In recent work, we have demonstrated that the image plane does not have to be focused on the micro-lenses of the plenoptic camera in order to capture the angular component of the radiance; the image could be formed any distance from the micro-lenses, and the angular information could still be derived from the individual micro-lens images (Lumsdaine, Georgiev 2008). But this type of focusing made apparent an additional problem: how to focus the micro-lenses precisely so that we extract the maximum resolution from the micro-lens images.

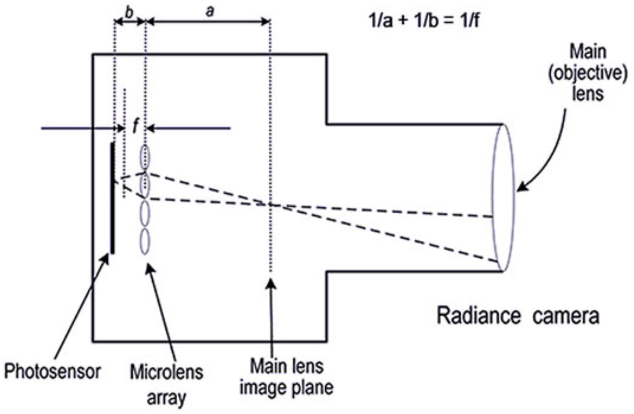

Figure 4. Our focused plenoptic camera with Keplerian imaging.

As it turns out, the solution is based on the two types of telescope design, Galilean and Keplerian. Essentially, the telescopic function is materialized inside the camera. The image plane focused in front of the micro-lenses acts like a scene one might view through a telescope eyepiece; the combination of the telescope eyepiece and the lens of the eye is replaced by each micro-lens; and the CCD array acts like the retina of the eye. The little micro-images created by the micro-lenses are potentially differently inverted and resized “puzzle pieces” derived from the full image. The structure of this “integral” image depends on depth at each of the above micro-images. Our algorithm resizes, inverts, and appropriately patches and blends them together in order to create the final rendering of the image.

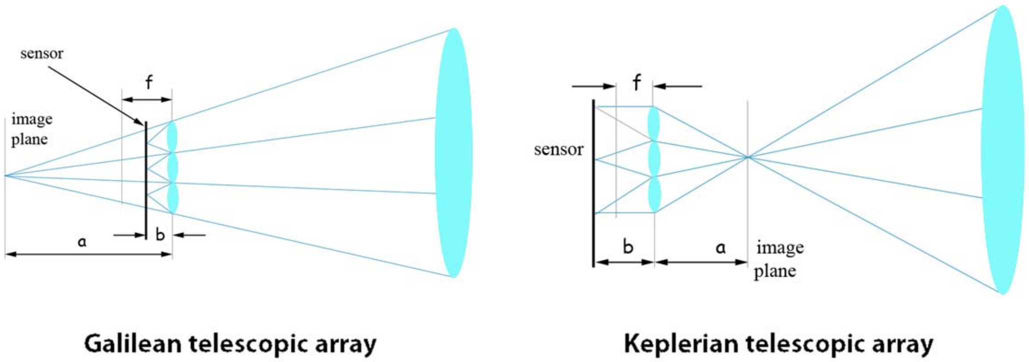

*Figure 5. Plenoptic Camera 2.0: Galilean and Keplerian imaging.**(Insert footnote 2)*With the “Plenoptic 2.0” camera design, the micro-lenses must be shifted exactly the right distance from the sensor so that

Deleted: ¶

Deleted: this

they project a correctly focused image on it. Figure 5 represents the optics of the two versions of our Plenoptic 2.0 camera, which not only has higher resolution but also performs sampling in a more flexible way than earlier plenoptic cameras. (Insert footnote 3) By spacing the micro-lenses from the sensor, we can fine-tune the slope and size (extent) of the area of optical phase space sampled by each micro-lens (Figure 6).

Deleted: “

Deleted: ”

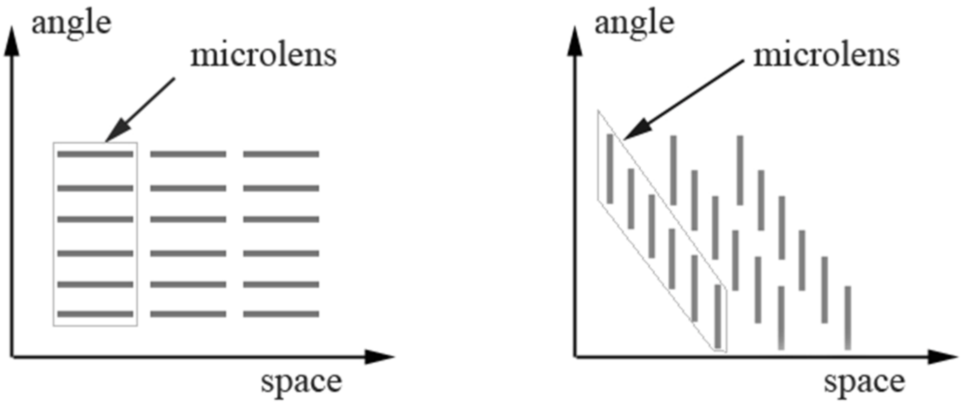

Plenoptic

Plenoptic 2.0

Figure 6. Sampling in Plenoptic 1.0 and Plenoptic 2.0 cameras.The figure on the left represents traditional plenoptic sampling. The figure on the right represents the sampling pattern in the Plenoptic 2.0 camera in which the redundancy of angular information allows for an increase in spatial resolution. Both left and right images include three micro-lenses positioned in optical phase space. Each micro-lens is a group of 6 pixels (gray bars). On the left, these bars are horizontal (covering a lot of space). On the right, the bars are vertical, covering little space, creating high resolution.

By adjusting the combination of the angle and the size of the sampled patch of optical phase space in each micro-lens (Figure 6), we can precisely select the spatio-angular trade-off point. In fact, we are able to reach a very low angular resolution not practical with prior plenoptic cameras because edge effects would generate noise. (Insert footnote 4) Thus, a much higher spatial resolution is possible while more flexibility is enabled in the choice of a spatial resolution trade-off.In computational photography, hardware and software must advance together. Therefore, with our Plenoptic 2.0 camera design, we have also developed new algorithms for rendering the final image. Before our improvement in the resolution, the plenoptic camera captured angular information but captured low spatial resolution (300x300 pixel images). With a Plenoptic 2.0 camera with large microlenses (more than 3,000 pixels behind each lens), we have been able to achieve up to a 500x increase in spatial resolution of the final image in terms of the number of pixels, at the expense of angular resolution (Figure 7). An improvement of 10x to 20x might be more practical for a real-world camera. Our algorithms create views from different vantage points, refocus the image, and enable other adjustments, all at high resolution. Earlier plenoptic cameras captured small images, typically 0.1 megapixels. With the new camera it will be practical to capture final images of around 2 megapixels based on a 39 megapixel sensor, and this can be improved further. We believe this represents a breakthrough in the practical applicability of the integral/plenoptic family of camera designs.

Deleted: ¶

Deleted: fine-tuned

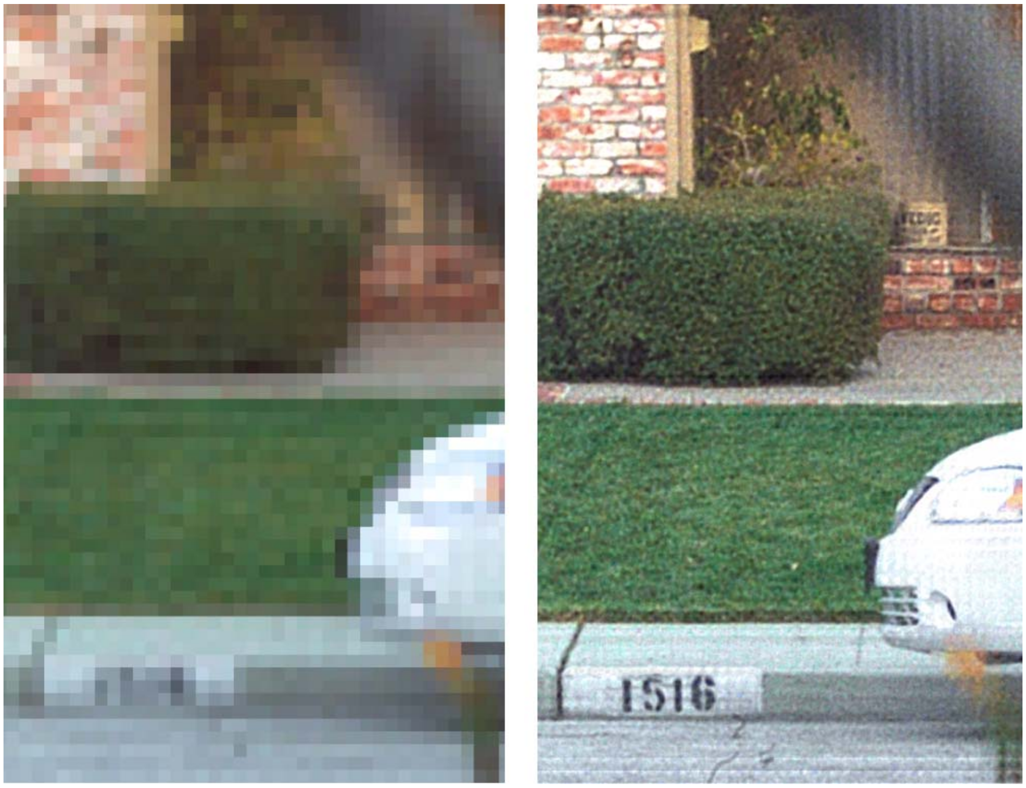

Figure. 7: Resolution of Plenoptic 1.0 and Plenoptic 2.0 cameras. Note: we show only 5% of the image.

What does this innovation enable?

For 100 years, Lippmann's insights have been within the realm of possibility. But trials and research have not yet produced a solution practical enough for our customers. With the current generation of computer power and camera hardware, these ideas are finally becoming practical.

But what exactly do these innovations mean in the realm of digital photography? Image editing software like Adobe Photoshop enables a new level of manipulation of the data collected in computational photography. For example, Adobe Labs has been working on a technique for extending depth of field by combining the focus information in multiple images in Photoshop. This research is directly applicable to the new plenoptic camera. But with integral photography, we go even further by capturing much more of the radiance light field information in a single shot. In addition to virtually refocusing an image after the fact, many other image manipulations become possible, including:

- The aperture and depth of focus can be adjusted **after** the image is captured.
- The image can be refocused with pixel-by-pixel control; paint any part of an image into or out of focus with a "Focus Brush" in a way that is impossible with a camera and lens alone.
- The vantage point of the camera can be adjusted, effectively shifting the position of the camera **after** the image is captured.
- Objects in a photograph can be selected automatically based on depth information rather than based on imperfect color recognition methods now used.

- Objects can be inserted into a scene, with proper occlusion based on depth.
- All of these techniques are possible with video as well, in software like Adobe After Effects.

These features allow photographers to defer image adjustments to the digital darkroom, freeing them to spend their time behind the camera pursuing the art of their craft: capturing the scene, the object, the moment. They also make possible entirely new image editing capabilities, so that we can do more than a photographer could possibly do in the field, with previously unimaginable control of the three-dimensional nature of images. Photographers now can create the physically impossible image!

Even as Adobe promotes this innovative research into the future of photography, we are always building on the shoulders of giants. We referred to the work of greats Galileo and Kepler, as well as more recent minds such as Levoy and Ng. And of course we began this paper with the powerful, imaginative work of Lippmann in 1908. His ideas have taken 100 years to reach a practical state – we can only imagine what the next 100 years will bring.

Footnotes

Formatted: Font: Bold

1a. Technology sets a limit to how small pixels can be, thus limiting spatial resolution achievable at the image sensor. Diffraction sets an additional limit to resolution. The lesser of these two limits gives us the maximum number of image pixels available in an image. It is the product of the number of spatial pixels times the number of angular pixels at each image location. In theory, camera optics can be designed so that the same product is derived from different spatial and angular arrangements, such as 100=10x10 or 100=50x2.

1b. In traditional plenoptic camera designs micro-lenses must work at the same numerical aperture as the main lens. If micro-lenses work as diffraction limited, their sizes will be much bigger than the diffraction limited resolution of the main lens. Each micro-lens captures “multiple pixels.” This results in aliasing. Placing a diffuser in front of the micro-lenses to reduce aliasing can reduce angular resolution, so the problem of aliasing makes these cameras difficult to build at high quality. Our Plenoptic 2.0 designs use a relay system (explained below) to capture spatial and angular information at the diffraction limit without aliasing. This design is able to achieve the maximum resolution optically possible in terms of the total number of pixels.

2. In Figure 5, when “a” becomes longer, angular resolution becomes higher, and spatial resolution becomes lower.

Formatted: Font: (Default) Arial, 10 pt, Font color: Auto

3. Using “Plenoptic 1.0” to refer to earlier versions of the Plenoptic camera, we can identify two primary differences between Plenoptic 1.0 and 2.0 cameras: A) Plenoptic 1.0 puts the image plane of the main camera lens on the micro-lens array; Plenoptic 2.0 focuses the image of the main lens at distance “a” from the micro-lenses; B) Plenoptic 1.0 places the micro-lenses at a distance “f” from the sensor, i.e. focused at optical infinity; Plenoptic 2.0 places the micro-lenses at distance “b” from the sensor focused at distance “a” from the micro-lens – exactly where the image of the main lens is formed. Thus, in the Plenoptic 2.0, micro-lenses are exactly focused.

4. Traditional plenoptic cameras achieve high spatial resolution by increasing the number of micro-lenses. Each micro-lens creates a small image of a few pixels, some of which are edge pixels which are noisy because they are not completely covered by the image. Edge pixel coverage is random in that we cannot determine exactly how much of the pixel is covered. Since smaller images have a higher percentage of edge pixels, increasing the number of micro-lenses increases the number of edge effect pixels, or the amount of noise, in the image. In the Plenoptic 2.0 camera, high spatial resolution is not achieved with

small images of the individual micro-lenses, but with a high slope of each micro-lens strip (Figure 6 right). The image is large and thus it has few edge pixels and little noise.

Deleted: ¶
¶

Formatted: Font: (Default) Calibri, 11 pt

References

Adelson, E.H., and Wang, J.Y.A. 1992. Single Lens Stereo with a Plenoptic Camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 14, No. 2, February.

Georgiev, T., Zheng, K. C., Curless, B., Salesin, D., Nayar, S., and Intwala, C. 2006. Spatio-angular Resolution Trade-off in Integral Photography. *Eurographics Symposium on Rendering*.

Levoy, M., and Hanrahan, P. 1996. Light Field Rendering. *ACM Trans. Graph*., 31-42.

Lippmann, G. 1908. Epreuves reversibles donnant la sensation du relief. *Journal of Physics* 7, 4, 821-825.

Lumsdaine, A., and Georgiev, T. 2008. *Adobe Technical Report*.

Ng, R., Levoy, M., Brledif, M., Duval, G., Horowitz, M., Hanrahan, P. 2005. Light Field Photography with a Hand-Held Plenoptic Camera. *Stanford University Computer Science Tech Report CSTR 2005-02*.