Superresolution with Plenoptic Camera 2.0

Todor Georgiev and Andrew Lumsdaine

Figure 1: *A bird on the antenna. The left image is rendered from our radiance data with the earlier plenoptic 2.0 algorithm. The right image is a superresolved rendering from the same data.*

Abstract

This work is based on the plenoptic 2.0 camera, which captures an array of real images focused on the object. We show that this very fact makes it possible to use the camera data with super-resolution techniques, which enables the focused plenoptic camera to achieve high spatial resolution. We derive the conditions under which the focused plenoptic camera can capture radiance data suitable for super resolution. We develop an algorithm for super resolving those images. Experimental results are presented that show a 9× increase in spatial resolution compared to the basic plenoptic 2.0 rendering approach.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision, Imaging Geometry, Super Resolution]:

1. Introduction

In 1908, Lippmann introduced his idea of using a microlens array to capture the light field of a scene and produce what he called *integral photographs* [Lip08b]. These photographs captured not only a 2D picture, but also "the 3D relief" of the scene [Lip08a]. In 1928, Ives subsequently added an objective lens to Lippmann's microlens array [Ive28]. Since then, many other researchers and scientists have continued to advance integral photography.

With the advent of digital photography, significant new opportunities to investigate integral photography became available. In 1992, Adelson introduced the "plenoptic camera," a digital version of Lippmann's microlens array designed to solve problems in computer vision [AW92]. The *light field*, introduced to the computer graphics community in 1996, became a framework for analyzing and processing radiance data [LH96]. In 2005, Ng improved the plenoptic camera and introduced new digital-processing methods, including refocusing [NLB*05, Ng06].

Capturing image data with the plenoptic camera makes possible greater processing capabilities and solves many of the problems faced by photographers using conventional

digital cameras. Unfortunately, traditional plenoptic cameras render images at very low resolution. For example, images rendered from Ng's camera data have a final resolution of 300×300 pixels. The early plenoptic cameras had such low resolution because of the way they sampled the 4D radiance of the scene. This sampling was based on the assumption that spatial and angular information was captured independently. The position of the microlenses captured spatial information while the pixels under each microlens sampled the angular distribution at the position of each microlens.

A different approach (called the *plenoptic 2.0 camera*, or the *focused plenoptic camera*) can render final images at much higher resolution than the original plenoptic camera. A number of researchers have noted this camera as an interesting alternative to the plenoptic camera. Surprisingly, a figure describing the idea can be seen in the early 1908 work of Lippmann [Lip08b], but this camera has been independently considered by Ng [Ng06], Fife [FEW08], and Lumsdaine [LG08]. The camera is structurally different from the earlier *plenoptic 1.0 camera* with respect to microlens placement, microlens focus, and, most importantly, with respect to assumptions made about sampling positional information. Lens placement in this new camera focuses the microlenses on the image inside the camera rather than at infinity. This structural change introduces a flexible trade-off in the sampling of spatial and angular dimensions and allows multiple pixels from each microlens to be rendered into the final image.

In the existing approaches to plenoptic 2.0 rendering, the trade-off between spatial and angular resolution is determined by parameters in the camera optics. Depending on depth, some portions of the scene may have more angular samples (and correspondingly, a lower spatial resolution) than others. As we will see in the paper, these additional samples of the same spatial area of the scene can be used to further increase the spatial resolution of that area through the use of super resolution.

Super resolution is a widely used technique for increasing the resolution of images [NB03, BK02, BS98, EF97, Hun95, LS04, Sch02]. There are many different methods for employing super resolution, but they all rely on extracting subpixel information from multiple images of a given scene to produce one higher-resolution image. For lightfield capture based on arrays of separate cameras, such as the Stanford array [WJV*05], it seems clear that super resolution would be directly applicable.

As we will show in the paper, the plenoptic 2.0 camera works as an array of cameras. These cameras are focused on the photographed object, a unique feature that distinguishes the plenoptic 2.0 camera from the conventional plenoptic (1.0) camera. Based on this feature, we will develop and apply super-resolution techniques to the rendering of light fields captured by the plenoptic 2.0 camera.

Concurrent with the initial submission of this paper to EGSR 2009, the paper [BZF09] appeared in publication. The approach presented in [BZF09] applies superresolution to the sub-sampled images obtained at each of the different views captured accross the microlens array. The authors develop a sophisticated image formation and restoration model and apply a blind deconvolution restoration approach. Although there has been only a very brief window of time for us to study this related work, it appears to be of high quality and great value. There are several key differences with our approach. Our work specifically interprets the plenoptic camera in terms of the focused plenoptic camera. This results in the use a different image formation model, one in which we account explicitly for sub-pixel registration. Based on this model, we present a camera design that is guaranteed to produce sub-pixel offsets between microlens images for objects at optical infinity. Those offsets are precisely computed from camera geometry and very reliable. Even if our approach is specifically targeted at super resolving for objects at infinity, we provide an analysis of the opportunities for superresolution at different depths in the scene and discuss the interactions between magnification factor and superresolution. Finally, our approach uses an experimentallay measured kernel in conjunction with the sub-pixel registration.

The plan of the paper is as follows: In Section 2, we will compare and contrast the plenoptic 1.0 and plenoptic 2.0 cameras, particularly in terms of trade-offs in spatial and angular resolution. In addition, we will show that the plenoptic 2.0 camera is equivalent to an array of cameras, each one focused on the object being imaged. Section 3 will describe our camera setting and our approach for applying super resolution. Experimental results will be presented in Section 4 and Section 5 will conclude.

Contributions:

1. We find the expressions for the positions and parameters at which super resolution is possible with the plenoptic 2.0 camera.
2. We develop and analyze several classes of super-resolution algorithms for the plenoptic 2.0 camera.
3. We propose and demonstrate a practical method of super resolving objects at optical infinity with a plenoptic 2.0 camera.

2. Plenoptic Cameras

2.1. The Plenoptic Camera 1.0

A traditional plenoptic camera [AW92, NLB*05, Ng06] consists of a main lens and a microlens array placed at distance f in front of a sensor (see Figure 2). Microlenses have aperture d and focal length f and, are assumed to be equally spaced at interval d. The main lens of the camera is assumed to be focused at the microlens plane and the microlenses are focused at infinity.

Considering that the focal length of the main camera lens

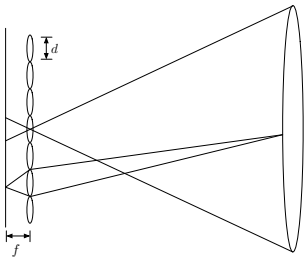

Figure 2: *Plenoptic 1.0 (traditional) camera. The main lens is focused at the microlens plane and the microlenses are focused at optical infinity (equivalently, the main lens).*

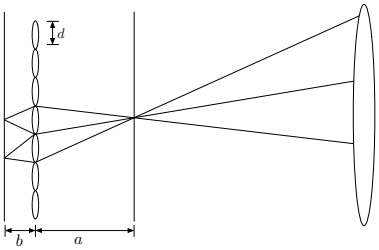

Figure 3: *Plenoptic 2.0 (focused) camera. The microlenses image the focal plane of the main lens.*

is much greater than the focal length of the microlenses, we have each microcamera focused at the main camera lens aperture, and not on the object being photographed. Each microlens image is thus completely defocused relative to that object, and it represents only the angular distribution of the radiance. As a rule, these microimages look blurry and do not represent a human-recognizable image. Since each microlens image samples a given location depending on its position and spans the same angular range, rendering an image from a plenoptic 1.0 radiance can be accomplished by simply integrating all of the pixels under each microlens. Integrating a fixed small portion of the pixels under each microlens generates an image of one certain view. In all cases a microlens contributes only to a single pixel in the final image.

2.2. The Plenoptic 2.0 Camera

An alternative version of the plenoptic camera can be constructed from an array of microcameras focused on the image plane of the main camera lens instead of at infinity (see Figure 3). With this system, each microcamera is reimaging the main lens image onto the sensor. We will refer to this version as *plenoptic 2.0 camera*. With this camera, the microlenses form an array of true images of the main lens image as a relay system. If the main lens forms an image behind the microlenses, it would still be possible to focus them on that virtual image so that they form a real image on the sensor. In both cases, the microlens imaging is described by the lens equation $1/a + 1/b = 1/f$, with, respectively, positive or negative a. When remapped onto the sensor, the image of the main lens is reduced in size. We denote this reduction as $m = a/b$.

As a result of this scaling, the spatial resolution of the radiance captured by the plenoptic camera is a function of the resolution of the microlens images and the amount of overlap in rendering, not of the number of microlenses. This *decoupling* of resolution and number of microlenses is a critical observation that distinguishes the focused from the traditional plenoptic camera.

Another difference between the plenoptic 1.0 camera and the plenoptic 2.0 camera is in the nature of the information that is captured by each microlens. In the plenoptic 1.0 camera, each microlens images one position in the scene, capturing all of the angular information there. In the plenoptic 2.0 camera, different microlenses capture the same position; angular information is spread across microlenses. Accordingly, to render plenoptic 2.0 images, we must integrate *across* microlens images, rather than within a single microlens image. That is, assuming we are “imaging the image” that is in focus, we integrate the points in the microlenses that correspond to the same position in the image by overlapping them at a fixed pitch.

In order to apply super-resolution techniques, we need to precisely characterize the microcamera array. In particular, the array of microlenses together with the main lens is equivalent to an array of cameras due to this system’s relay imaging mode of work (see Figure 3). An array of microcameras observe the "object" in front of them. This "object" is the aerial 3D image of the scene, placed behind the main camera lens. Since super resolution is applicable to an array of cameras imaging an object, it is applicable to plenoptic 2.0 imaging.

3. Super Resolution for Plenoptic 2.0

3.1. Super Resolution Model

The super-resolution problem is to recover a high-resolution source from multiple low-resolution observations. The low-resolution observations may be produced in a variety of different ways, depending on the application. They may be captured by a camera array, a single shifted camera, or they may even be different frames of a video sequence.

The image-acquisition process in the focused plenoptic camera is shown in Figure 4 and is modeled as follows. A pixel p_1 under microlens 1 samples radiance within a certain angle from a given spatial area in the main lens image (in front of the microlenses). In the same way, a pixel p_2 under microlens 2 samples an area partially overlapping with the area sampled by p_1, i.e., an area shifted by a sub-pixel

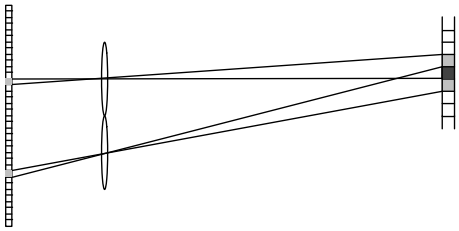

Figure 4: *Low resolution acquisition of a high-resolution image. Microlenses sample overlapping regions of the high-resolution image generated by the main camera lens.*

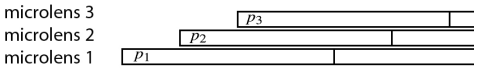

Figure 5: *Overlapping pixels (p_1, p_2, p_3) in the same sampling area. Pixels are mapped from the sensor to the area sampled and placed on top of each other in space.*

amount (See Figure 5). A pixel p_3 under microlens 3 samples an area partially overlapping with the area sampled by p_1 and p_2. And so on.

Each of those pixels samples a version of the outside world scene, blurred through the kernel of the camera optics. This includes both the main lens and microlens corresponding to the pixel. Also, the final pixel value is the result of the convolution of that blurred image with the point-spread function of the pixel sensor's responsivity. The total kernel is represented as H with an added noise term. This is the typical analysis of super resolution, now adapted to the focused plenoptic camera:

$$b = Hx + n. \tag{1}$$

Here, b represents the collected low-resolution observed images, H is the blur matrix, n is a noise term, and x is the high-resolution image we wish to recover.

Recovering x is then cast as a minimization problem:

$$\min_x \left\{ \|Hx - b\|_2^2 + \alpha R(x) \right\}, \tag{2}$$

where $R(f)$ is a regularization term whose choice depends on the application and desired solution characteristics. Formulating and solving this problem efficiently and effectively in different application areas is an active area of research [NB03, BK02, LS04, Sch02].

Key to the success of any super-resolution approach is that there be nonintegral (subpixel) shifts between different aliased observations of the high-resolution images. In the general case, estimating these shifts (and, consequently, forming H) is also part of the super-resolution problem. In

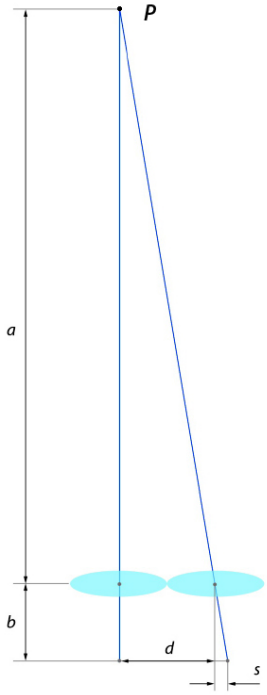

Figure 6: *Geometry of our data capture for super resolution with plenoptic 2.0 camera.*

the case of our focused plenoptic camera, some of the super-resolution problem is simplified as we have an array of cameras spaced with predetermined micron precision. On the other hand, the shift between features from one microlens image to the next depends on the scene and the camera optics.

3.2. Plenoptic 2.0 Camera Design for Super Resolution

In this section, we consider how to design the optics of the plenoptic 2.0 camera to best support super resolution.

The camera can be viewed as a relay imaging system, an array of microcameras focused on the image created by the main camera lens. Consider one luminous point P in the main lens image. Figure 6 represents the imaging of this point in two microcameras. To be amenable to super resolution, we want the distance between the images of the point to be a nonintegral pixel value.

In the figure, we see that $d/a = s/b$. From this, we derive $a = db/s$. Since, in general, the distance d between microlens centers is already not an integer, it would not be appropriate to constrain s as a way of controlling the subpixel shift. Rather, let the next integer larger than d be $\Delta = d + x$ and let $s = x + t$. Since we know the pixel size and d with great precision, we know x. Then, t is the translation from

the integer pixel location to the image of the observed point. In this way, $t = s - x$ is the quantity that we require to have a nonintegral value.

Note that there are multiple regions in the scene (multiple values of a and b) for which t will have a nonintegral value. For instance, we can take t to be 0.5 pixels, but we could also take it to be 1.5, or 2.5, or, in general, 0.5 + n for $n = 0, 1, 2, 3, \ldots$. After super resolving, these provide the same $2\times$ increase in the resolution.

The general case is $t = k + n$, where k is a fraction less than 1. Different types of super resolution can be designed with different k. With this notation, our general equation can be written as

$$a = \frac{db}{x + k + n}. \tag{3}$$

Super resolution is achieved with microimages shifted by $\Delta + k + n$ pixels.

In the plenoptic 2.0 camera, the portion of the scene that is at optical infinity (i.e., imaged at the largest distance from the microlenses) will have the greatest reduction in size. That is, the lowest spatial resolution under plenoptic 2.0 rendering. At the same time, since it is the farthest from the microlenses, it has the most angular samples. The low resolution and the availability of the most angular samples also means that this region of the scene is the most important to use for support of super resolution.

Different depths in the scene are imaged at different distances a in front of the microlenses. This creates a difficulty for support of super resolution because the depths would super resolve at different values of k. Solving this problem requires subpixel registration among all microimages, which may be a difficult problem: The solution may be computationally expensive or unreliable for automatic super resolution. What's more, certain depths would not be super resolvable at all if the shift between microimages happens to be close or equal to an integral number of pixels. This type of problem has plagued conventional super-resolution methods for years, and they still remain too unreliable for such commercial image-processing products as Adobe Photoshop, for example.

However, the plenoptic 2.0 camera has a unique characteristic. There is one special depth in the scene, the depth of infinity, that is always mapped to the same location in front of the microlenses, one focal length from the main camera lens. Infinity is also the depth that benefits most from plenoptic 2.0 super resolution. This is also the depth that can be handled with highest precision for super resolution since it is fixed and subpixel correspondence is set and exactly known in advance.

For a given type of super resolution (defined by the fraction k) there are a number of planes that satisfy the subpixel shift condition. As they approach the microlens array those planes become denser and denser; at a certain point their position becomes hard to determine and unreliable (see Figure 7). The plane corresponding to infinity is the farthest from the microlens array at which there is image to capture. The separation between it and the previous one is the largest. This makes it the plane with most reliable correspondence, best for super resolution.

We design the camera such that infinity is super resolved directly, with registration provided by the camera geometry and the microlens pitch. This way we avoid estimation of registration from the imagery. At the same time our registration is much more precise and reliable.

3.3. Specific Design Examples

The parameters of our physical plenoptic camera are as follows. The microlens pitch is 500 μm and the sensor pixels are 6.8 μm. Thus, $d = 73.5294$ pixels, $\Delta = 74$ pixels, and $x = 0.4706$ pixels. The value for $b \approx 1.6$ mm could not be estimated with precision better than 0.1 mm because of the cover glass of the sensor. We estimated it approximately from known sensor parameters and independently from the microlens images at different F/numbers. In the end, we computed $db \approx 120mm$. Note that a and b are measured in millimeters while everything else is measured in pixels (i.e. no units).

3.3.1. 2×2 Super Resolution

Suppose we wish to super resolve a plenoptic 2.0 image increasing the size two times in each direction. For 2×2 super resolution, we need $t = 0.5 + n$, where $n = 0, 1, 2, 3, \ldots$ and $a = db/(x + 0.5 + n)$. With the parameters of our camera above, we have approximately $a \approx 120/(1 + n)$ measured in millimeters. The values of n at which the camera super resolves and the corresponding distances a (in millimeters) are given in the table below.

a	120	60	40	30	24	20	17.1	15	13.3	12	10.9	10	9.2
n	0	1	2	3	4	5	6	7	8	9	10	11	12

3.3.2. 3×3 Super Resolution

If we want to super resolve a plenoptic 2.0 image three times in each direction, we need $t = 1/3 + n$, where $n = 0, 1, 2, 3, \ldots$ and $a = db/(x + 1/3 + n)$. With the parameters of our camera above, we have approximately $a \approx 120/(0.8 + n)$ measured in millimeters.

a	150	66.6	42.8	31.6	25	20.7	17.6	15.4	13.6	12.2	11.1
n	0	1	2	3	4	5	6	7	8	9	10

The depth planes at which the above two types of super resolution work are represented in Figure 7.

Other types of super resolution, such as 5×5, and so on, can be designed easily.

The results in this paper are generated by selecting $n = 8$ in the table for 3×3 super resolution, corresponding to a

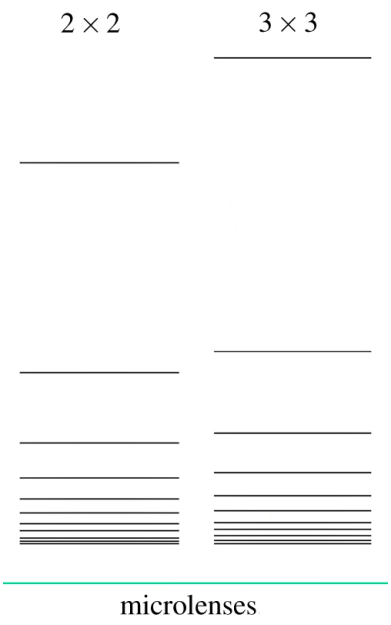

Figure 7: *Planes at different distances from the microlenses (represented in the vertical direction) at which plenoptic 2.0 camera super resolves in* 2×2 *and* 3×3 *mode.*

distance $a = 13.6$ mm in front of the microlenses. (We have chosen a relatively big value of n because for n lower than that the image is too far from the microlenses and too small.) We have moved the main camera lens forward with a 13-mm extension tube and fine tuned the position of the lens by changing the focus. Refocusing provides fine tuning by continuously moving the lens forward between 0 and 10 mm. Super resolution was performed at a shift of $\Delta + n = 82$ pixels between captured microimages.

4. Experimental Results

4.1. Camera

We are working with a medium format camera, using an 80-mm lens and a 39-megapixel digital back from Phase One. Pixel size is 6.8 μm. The lens is mounted on the camera with a 13-mm extension tube, which provides the needed spacing to establish an appropriate distance from the focal plane to the microlens array and the sensor. This setting is described in Section 3.2.

The microlens array is custom made by Leister Micro Optics. We have designed it to work with the sensor without removing the cover glass. For that purpose, the microlenses have focal length of 1.5 mm and the array is placed directly on the cover glass of the sensor, after removing the infrared filter. We have also crafted a way to provide variable additional spacing of up to 0.2 mm, which in our experience proved to be extremely helpful for fine tuning the microlens focus.

The pitch of the microlenses is 500 μm with a precision better than 1 μm. This precision makes subpixel registration and our method of super resolution possible. The large pitch makes our microlenses work with an F-number of the main lens as low as 3. Note that if such an array were used with a traditional plenoptic camera, it would result in a final image of size less than 100×100 pixels after rendering. That is why all previous plenoptic camera designs required removal of the cover glass and a very small distance to the sensor—in order to accommodate a low F-number at small pitch. Our design is possible only because of the plenoptic 2.0 approach we are taking.

Next we want to estimate the range of depths in the real world at which our super resolution works. As noted in the previous section, the image of infinity is formed at distance $13.6mm$ from the microlenses, corresponding to $n = 8$. The next closer plane good for 3×3 super resolution would be at $12.2mm$, and between them there is a plane where super resolution would fail. Let's assume that our image is well super resolved within $0.5mm$ (from $13.1mm$ to $13.6mm$). Consider the lens equation for the main camera lens:

$$(A - F)(B - F) = F^2 \tag{4}$$

where $F = 80mm$ is the focal length of the main lens, A is the distance to the object, and B is the distance to the image. Our estimate above that $B - F = 0.5mm$ leads to a distance $A = 12.8m$. Anything that is located at more than $13m$ from the camera is well super resolved.

Note that we can bring this plane closer based on selecting different camera parameters.

4.2. Algorithm

To solve equation (1) in our case, we use the following approach.

1. Create a high-resolution observed image b by interleaving pixels from adjacent microlens images. For the experiments shown here, we use a 3×3 resolution increase, so each microlens image interleaves pixels from its eight nearest neighbors.
2. Solve equation (1) with an appropriate computational method. For the results shown here, we used the approach and software described in [LFDF07, LFDF], with Gaussian and sparse priors. The kernel used for deconvolution was obtained by imaging a point light source (pinhole).

4.3. Images

Figure 8 shows a portion of a lightfield/radiance captured by our camera. No part of the image is in focus at the im-

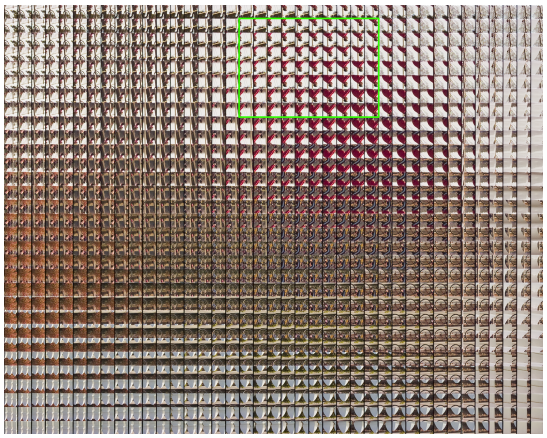

Figure 8: *A portion of the radiance recorded by our camera.*

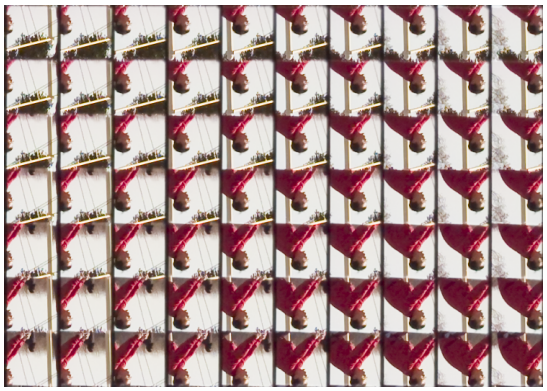

Figure 9: *Extreme close up of a portion of the green rectangle in Figure 8. The microimages are inverted and well-focused. The repeated image of an individual riding a bike can be seen.*

age plane, hence the lightfield image appears blurry at a macro level. However, if we closely examine the microimages themselves, we see a different story. They are well focused, as shown in Figure 9, which is a zoom-in into the green rectangle. Note that for efficient use of sensor space we have installed a square main lens aperture, so our microimages are squares and not circles.

Figure 10 shows a stereo rendering of the radiance rendered using the plenoptic 2.0 algorithm (without super resolution). To see the effects of different rendering approaches more clearly, we consider smaller portions of the image. Traditional lightfield rendering, with one pixel per microlens, yields an image with very low resolution, as shown in Figure 11. Alternatively, the plenoptic 2.0 (full resolution) rendering algorithm enables significant resolution improvement, as shown in Figure 12. A different view of the same data rendered with the plenoptic 2.0 super resolution algorithm (as described in Section 4.2) is shown in Figure 13.

Figure 14 shows a close up of the front wheel of the bicycle when rendered with the plenoptic 2.0 algorithm—pixelation is obvious. Figure 15 shows a close up of the front wheel of the bicycle when rendered with plenoptic 2.0 and super resolution, using the algorithm described in Section 4.2.

Another example is the image of a bird on the antenna (Figure 1), which is a crop from a much bigger picture. If we zoom in, we can clearly see pixelation in the plenoptic 2.0 rendering and the improvement gained with our super resolution rendering.

The complete files of those images, as well as the input lightfields we used, are available with the electronic version of this paper.

5. Conclusion

The radiance (or plenoptic function) carries a significant amount of information. This information can be used to generate novel effects when rendering. With the plenoptic 2.0 camera, we are able to make deliberate spatio-angular tradeoffs and obtain significant improvements in spatial resolution. With the application of super-resolution techniques as presented in this paper, we are able to push the attainable spatial resolution even further. One factor that has limited the adoption of plenoptic cameras until now has been the relatively low available resolution. The combination of the focused plenoptic camera with super resolution enables images of sizes acceptable to modern photographers, making lightfield photography immediately practical.

References

[AW92] ADELSON T., WANG J.: Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1992), 99–106.

[BK02] BAKER S., KANADE T.: Limits on super-resolution and how to break them. *IEEE Transactions on Pattern Analysis and Machine ...* (January 2002).

[BS98] BORMAN S., STEVENSON R.: Super-resolution from image sequences-a review. *Proceedings of the 1998 Midwest Symposium on Circuits and ...* (January 1998).

[BZF09] BISHOP T. E., ZANETTI S., FAVARO P.: Light field superresolution. In *International Conference on Computational Photography* (2009).

[EF97] ELAD M., FEUER A.: Restoration of a single superresolution image from several blurred, noisy, and undersampled measured .... *Image Processing* (January 1997).

[FEW08] FIFE K., EL GAMAL A., WONG H.-S. P.: A 3MPixel multi-aperture image sensor with 0.7 μm pixels

Figure 10: *A stereo pair rendered from our captured data (Figure 8) with the plenoptic 2.0 rendering algorithm. Intended for crossed eye viewing (left and right image switched).*

Figure 12: *A view rendered from the same data with the plenoptic 2.0 (full resolution) rendering algorithm.*

Figure 11: *Portion of the same image rendered with the trafitional (plenoptic 1.0) lightfield rendering algorithm.*

Figure 13: *A slightly different view rendered from the same data with our plenoptic 2.0 super resolution algorithm. It includes some of the white fence in the close foreground.*

Figure 14: *Close up of the front wheel of the bicycle, using plenoptic 2.0 rendering.*

Figure 15: *Close up of the front wheel of the bicycle, using plenoptic 2.0 rendering and super resolution.*

in 0.11 μm CMOS. In *IEEE ISSCC Digest of Technical Papers* (February 2008), pp. 48–49.

[Hun95] HUNT B.: Super-resolution of images: algorithms, principles, performance. *International Journal of Imaging Systems and Technology* (January 1995).

[Ive28] IVES H. E.: A camera for making parallax panoramagrams. *Journal of the Optical Society of America 17*, 4 (December 1928), 435–439.

[LFDF] LEVIN A., FERGUS R., DURAND F., FREEMAN W.: Sparse deblurring code. http://groups.csail.mit.edu/graphics/CodedAperture.

[LFDF07] LEVIN A., FERGUS R., DURAND F., FREEMAN W.: Image and depth from a conventional camera with a coded aperture. *ACM Transactions on Graphics, SIGGRAPH 2007 Conference Proceedings, San Diego, CA* (2007).

[LG08] LUMSDAINE A., GEORGIEV T.: *Full Resolution Lightfield Rendering*. Tech. rep., Adobe Systems, January 2008.

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. *Proceedings of the 23rd annual conference on Computer Graphics and Interactive Techniques* (January 1996).

[Lip08a] LIPPMANN G.: Épreuves réversibles donnant la sensation du relief. *Journal of Physics 7*, 4 (1908), 821–825.

[Lip08b] LIPPMANN G.: Épreuves réversibles. photographies intégrales. *Académie des sciences* (March 1908), 446–451.

[LS04] LIN Z., SHUM H.: Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE Transactions on Pattern Analysis and Machine Intelligence 26*, 1 (January 2004), 83–97.

[NB03] NG M. K., BOSE N. K.: Mathematical analysis of super-resolution methodology. *Signal Processing Magazine, IEEE 20*, 3 (2003), 62–74.

[Ng06] NG R.: *Digital light field photography*. PhD thesis, Stanford University, Stanford, CA, USA, 2006. Adviser: Patrick Hanrahan.

[NLB*05] NG R., LEVOY M., BRÉDIF M., DUVAL G., HOROWITZ M., HANRAHAN P.: *Light Field Photography with a Hand-Held Plenoptic Camera*. Tech. Rep. CSTR 2005-02, Stanford University Computer Science, Apr. 2005.

[Sch02] SCHULTZ R.: Super-resolution enhancement of native digital video versus digitized NTSC sequences. In *Proceedings of the Fifth IEEE Southwest Symposium on Image Analysis and Interpretation* (2002), pp. 193–197.

[WJV*05] WILBURN B., JOSHI N., VAISH V., TALVALA E., ANTUNEZ E., BARTH A., ADAMS A., LEVOY M., HOROWITZ M.: High performance imaging using large camera arrays. In *ACM Transactions on Graphics* (2005).