CS348b Project: Light Field Camera Simulation

Zahid Hossain[*] Adam Spilfogel Backer[†] Yanlin Chen[‡]

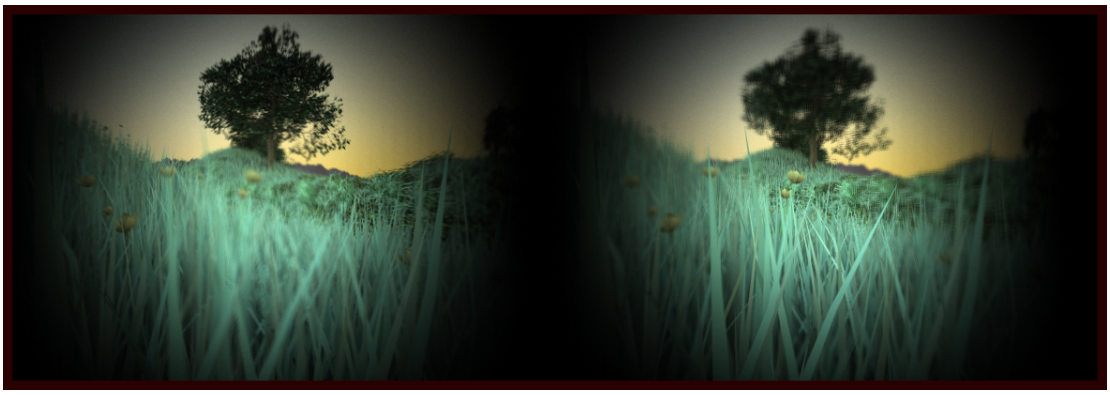

Figure 1: *Lightfield rendering to demonstrate focus at different depths. The left image is focused at a far plane near the tree while the right image is focused near flowers at a closer distance from the camera.*

Abstract

We report a simulation platform for a lightfield camera. Our methods may be applied to develop and refine future lightfield acquisition systems. Furthermore, the rendering pipeline of PBRT has been augmented to include our camera simulator, permitting the generation of refocus-able computer-generated artwork. We demonstrate our simulation platform on various imaging scenarios that may prove challenging for more conventional imaging models.

Keywords: lightfield, realistic camera simulation, imaging systems and applications

1 Introduction

Human beings are always looking forward to saving the very moment that we cherish. That's why we invented photography in the first place. And now we are even expecting a living picture that enables backward time travel. A plenoptic camera is then introduced [Ng et al. 2005], and we intend to bring it to the virtual graphics world. Such a camera, also referred to as light-field camera, captures all the power of light field, which contains the direction, color and brightness of the rays of light within the frame.

In this project, we simulate the realistic plenoptic camera inside pbrt system using a microlens array to capture 4D light field information and implement original refocusing algorithm to make plenoptic frames that can be refocused after they are 'taken. We have built multiple interesting test scenes with complex details to demonstrate the magic of a light field camera using pbrt.

[*]e-mail:zhossain@stanford.edu
[†]e-mai:abacker@stanford.edu
[‡]e-mail:yanlinc@stanford.edu

2 Lightfield Data and Synthetic Refocus

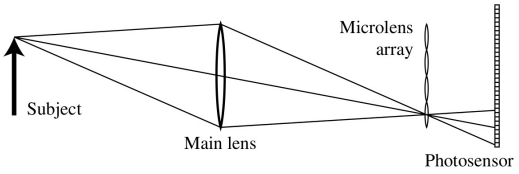

Figure 2: *A practical method for capturing lighfield by inserting a plane of microlens array between the photosensors and the main lens.*

A lightfield is a 4D function that captures the amount of light incident on a 2D spatial point (s, t) from an angular direction (u, v). A practical method to capture a lightfield is demonstrated in Figure 2 (image taken from Ng et al. [Ng et al. 2005]), where a conventional camera model is repurposed by inserting an extra plane of microlens array between the photosensor and the main lens. Unlike the conventional camera, where a single pixel records the integral of the incoming light cone, a pixel in lighfield camera records a single ray. In other words, each microlens deposits light ray onto a small patch of pixels – known as lenslet image – that are right behind it. This yields a raw image data that appears like Figure 3(a), where each small circular image patch corresponds to a single microlens. Mathematically, a single microlens is then a point-sample in the spatial dimensions (s, t) while each pixel in the lenslet image is a point-sample in the the angular (u, v) dimensions.

Since a microlens can see the entire main lens in front of it, the same u, v coordinates also span the entire main lens. Therefore, if take the same (u, v) pixel under every microlens and form an image that will be equivalent to capturing a pinhole image taken from the corresponding (u, v) position on the main lens. Such an image is called a subaperture image, and Figure 3(b) shows a montage of all subaperture images, which is also an alternative storage format for lightfield.

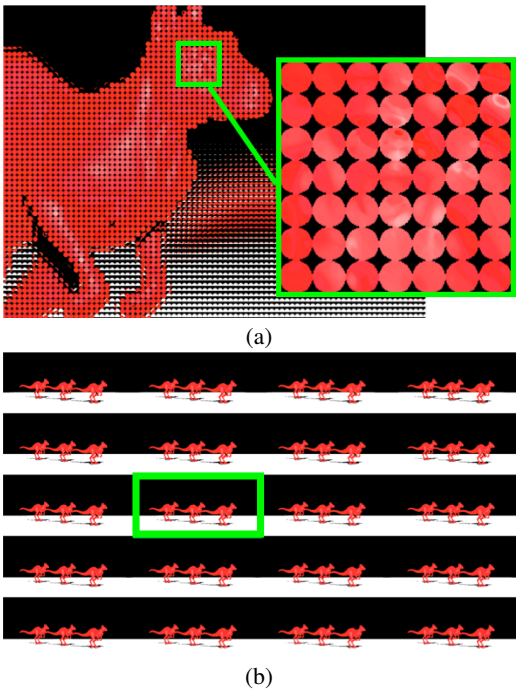

Figure 3: *(a) Raw lightfield data. Here each circular image patch (32×32 in this case) corresponds to a single microlens (b) Montage of subaperture images (one marked with a green rectangle). Each of these subaperture image was generated by sampling the same (u, v) lenslet pixel from every microlens. This format of data essentially is a simple rearrangement of the (a).*

2.1 Synthetic Refocusing

A lightfield function can be resampled appropriately to generate a synthetic focal plane as shown in Figure 4 (image taken from Ng et al. [Ng et al. 2005]). Rays are traced from the main lens (from u) that passes through a given synthetic pixel at s' to the captured lightfield at s. Now the goal is to to compute the intersection s given u and the distance of the synthetic plane from the main lens. Let this distance be parameterized by α, which is the fraction of the distance of the microlens array from the main lens. The intersection s is then given by:

$$s = \frac{1}{\alpha}\left(s' - (1-\alpha)u\right) \tag{1}$$

Note that we derived Equation 1 ourselves which is a much simplified version of Ng etal. [Ng et al. 2005]. Given u, and having computed s we can then lookup our recorded lightfield, $L(u, v, s, t)$ to integrate the irradiance at (s', t') which is given by the following synthetic refocus equation:

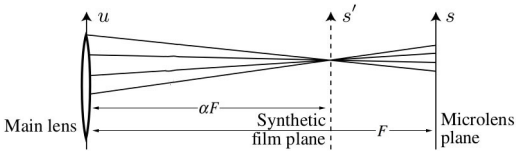

Figure 4: *Synthetic refocusing is performed at an arbitrary plane s' by tracing rays, through a given synthetic pixel, that pass through the subaperture spanned by u from the main lens to the captured lightfield plane s. The distance of the synthetic plane from the main lens is parameterized by α which is the fraction of total distance from the main lens to the microlens array.*

$$E(s', t') = \int_{u,v} L(u, v, \frac{1}{\alpha}\left(s' - (1-\alpha)u\right), \frac{1}{\alpha}\left(t' - (1-\alpha)v\right)\, dudv \tag{2}$$

Though the above equation is only a 2D integral we still performed a monte-carlo integration to mask off the subsampling aliasing with random noise that is visually more pleasant. We also performed a linear interpolation in all 4D dimensions to convert the sampled lightfield function into a C^0 continuous one.

Finally, we simply took off the $1/alpha$ term in Equation 1 to get rid of the scaling effect from our synthetically refocused image. This is equivalent to sampling the synthetic plane more densely as we move it closer to the main lens thus keeping the dimensions of the rendered image (in pixels) the same. Therefore, we used the following modified integral to compute our final refocused image:

$$E(s', t') = \int_{u,v} L(u, v, \left(s' - (1-\alpha)u\right), \left(t' - (1-\alpha)v\right)\, dudv \tag{3}$$

3 Results

3.1 Exploiting Lightfield Imaging

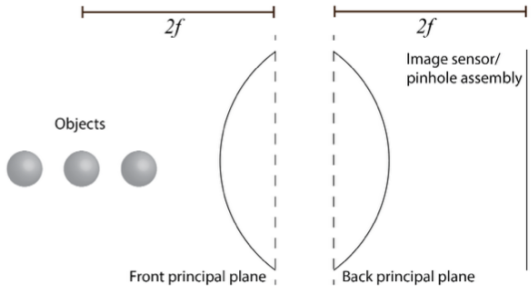

Figure 6: *Scene schematic for lightfield 'proving ground'*

In order to test our image processing algorithms and lightfield camera simulator, it was necessary for us to develop a contrived scene that was well suited for refocusing tasks. Figure 6 shows the layout for a scene that served as a 'proving ground' for our lightfield camera. The scene consisted of three objects– the middle object was placed a distance of two focal-lengths from the front principal plane of the main camera lens. The other two objects were placed slightly closer and further from the camera lens. The pinhole array and image sensor assembly was placed a distance two focal-lengths behind the back principal plane of the camera main lens. This configuration induced an approximately 1:1 magnification ratio between object and image points, and the proximity of the objects to the camera lens facilitated an imaging system with an exquisitely shallow depth-of-field. Our main lens consisted of a single bi-convex element composed of NBK7 glass (refractive index of 1.517). Using the lens maker's equation, we calculated the appropriate radius of curvature for the front and back surfaces of the lens. We For our scenes, we used a lens with an effective focal-length of 50 mm, a front/back surface radius of curvature of 48.76 mm, a thickness of 17.277 mm, and an aperture diameter of 20 mm. In Figure 7 we show a conventional image acquired using these camera specifications, and scene configurations.

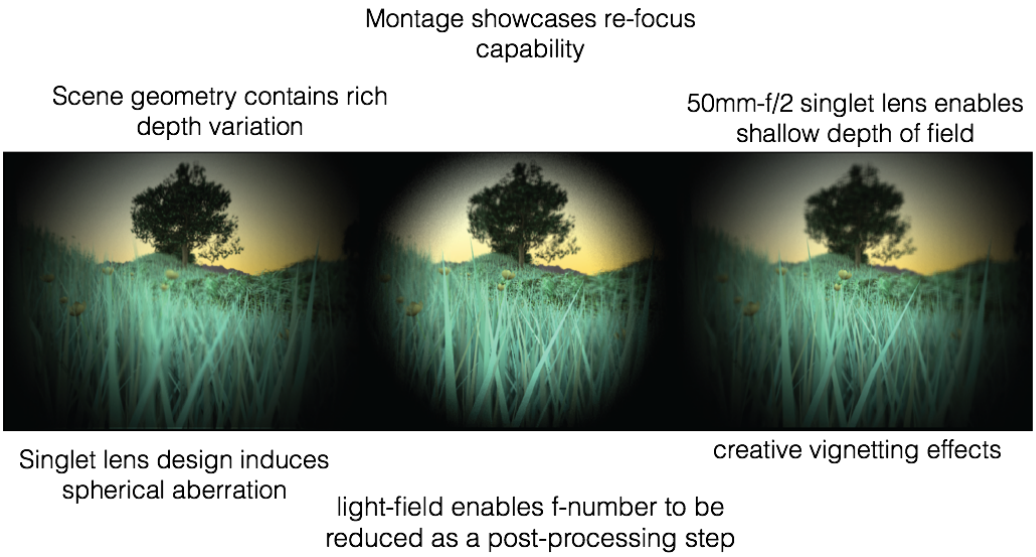

Figure 5: *Results of our main scene*

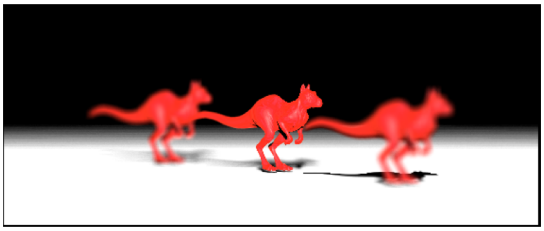

Figure 7: *Scene acquired without refocus applied. Only a single plane is in focus.*

To demonstrate basic refocusing capability, Figure 8 demonstrates our ability to move the plane of best-focus to points nearer or further from the camera main lens. In addition to changing the focus setting of a conventional camera, Lightfield data permits us to perform imaging tasks that would be mechanically impossible using a standard imaging system. For example, we augmented our refocusing algorithm such that different film-depths could be specified over different regions of the image sensor. In order to mimic the functions of a tilt-shift lens, we applied refocus using a gradient film depth (Figure 9.a). The result was an all-focused image (Figure 9.b). Furthermore, we manually defined regions on our image sensor to be refocused (Figure 9.c), thus permitting the near and far objects in our test scene to be refocused, while the central object was out of focus (Figure 9.d). Such an effect cannot be achieved using a conventional camera model.

Finally, in order to stress-test our refocusing algorithm, we tasked it with the following imaging scenario: We placed an object half-of one focal length from the front principal plane of our camera's main lens, and parked the image sensor half of one focal length from the back principal plane. In such a configuration, it is impossible for a conventional camera to focus on an object, since rays collected by the main lens will never converge (Figure 10.a). We display the data that would be recorded by a conventional camera using such a scene setup (Figure 10.b). A lightfield dataset of the scene was recorded, and the refocus algorithm was applied to it. Figure 10.c shows the resulting focused image. We term this effect 'impossible imaging'. The images obtained using this procedure

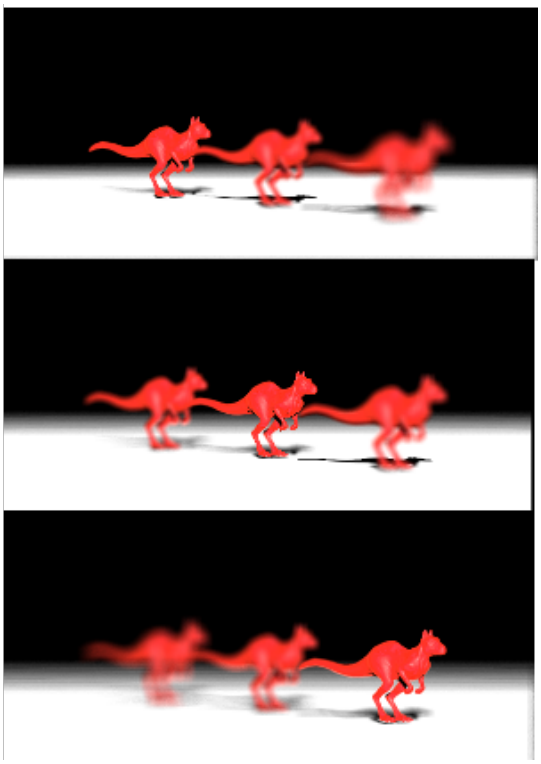

Figure 8: *Application of refocus to our 'proving ground' scene.*

were substantially blurred. We believe this effect is due to aliasing of the lightfield angular sampling, and may be a fruitful direction for future research.

3.2 Main Scene

We set up our main scene based on the resource `plant-dusk.pbrt` from [Deussen et al. 1998], which is a complex ecosystem model, rendered with an environment map

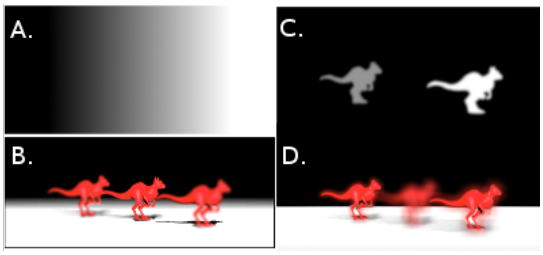

Figure 9: *Selective refocusing tasks applied to our 'proving ground' scene.*

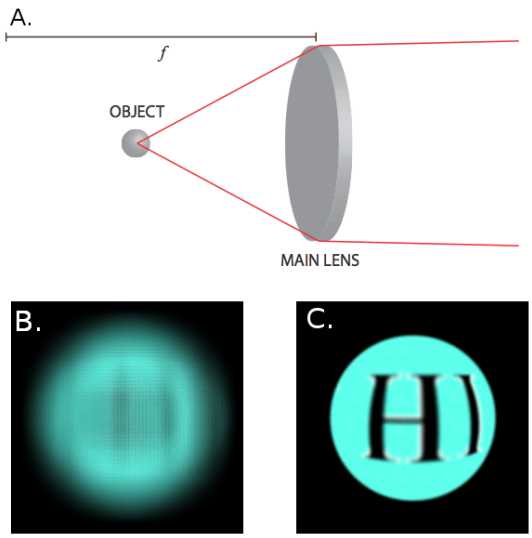

Figure 10: *'Impossible' imaging achieved with the lightfield camera.*

simulating illumination at dusk. Figure 5 shows our main scene for the rendering competition. The near plane is composed of lots of details of the grass, while the big tree is located far away.

Our plenoptic camera captures all the lightfield information in a single frame. Having applied our refocusing algorithm, one can easily (1) generate refocused image and (2) change the center of the camera after the picture is taken. The left one shown in Figure 5 is focused at the tree far away, the right near focused. The middle is a subaperture image taken from the center of the camera with a much smaller radius, which explains some noise around the edge of the lens.

During the post-processing after the image is taken, we wrote a script to first simulate the effect of changing depth of field, focusing from far to near, and then panning the camera up and down. We have made a video that demonstrate how we change the focus and move the camera. Notice the parallax effects as we pan the camera, due to the power of lightfield information.

4 Conclusion

We have successfully implemented a realistic lightfield camera inside pbrt system as well as sophisticated refocusing algorithm. We have tested our system on lots of scene that demonstrate the power of a lightfield camera including the ability of both refocusing and changing the viewpoints after a picture is taken.

Acknowledgements

We would like to sincerely thank Matt Pharr and Daniel Ritchie for teaching this wonderful course. Finally, we would also like to thank all the contributors and authors of PBRT along with Zach Devito for this life saving tool VDB (`https://github.com/zdevito/vdb`).

References

DEUSSEN, O., HANRAHAN, P., LINTERMANN, B., MĚCH, R., PHARR, M., AND PRUSINKIEWICZ, P. 1998. Realistic modeling and rendering of plant ecosystems. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM, 275–286.

NG, R., LEVOY, M., BRÉDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light field photography with a hand-held plenoptic camera. *Computer Science Technical Report CSTR 2*, 11.

A Extra Scenes

Figure 11: *San Miguel scene focused at different depth using a single spherical main lens. Note on the top row, how the foliage at the periphery of the main lens got focused earlier than the center part. This is caused by the spherical lens aberration that focuses light at the periphery at a closer distance than the actual focal length.*

a b c

Figure 12: *Comparison of various interpolation techniques. (a) Nearest Neighbor interpolation in all dimensions that produces the noisiest image (b) Linear interpolation in the spatial (s, t) dimensions that produces a little less noisier output. (c) Linear interpolation in all 4 dimensions that produces the best result.*