Wide-Baseline Light Fields using Ellipsoidal Mirrors

Michael De Zeeuw and Aswin C. Sankaranarayanan

Abstract—Traditional hand-held light field cameras only observe a small fraction of the cone of light emitted by a scene point. As a consequence, the study of interesting angular effects like iridescence are beyond the scope of such cameras. This paper envisions a new design for sensing light fields with wide baselines, so as to sense a significantly larger fraction of the cone of light emitted by scene points. Our system achieves this by imaging the scene, indirectly, through an ellipsoidal mirror. We show that an ellipsoidal mirror maps a wide cone of light from locations near one of its foci to a narrower cone at its other focus; thus, by placing a conventional light field camera at a focus, we can observe a wide-baseline light field from the scene near the other focus. We show via simulations and a lab prototype that wide-baseline light fields excel in the traditional applications involving changes in focus and perspective. Additionally, the larger cone of light that they observe allows the study of iridescence and thin-film interference. Perhaps surprisingly, the larger cone of light allows us to estimate surface normals of scene points by reasoning about their visibility.

Index Terms—Light fields, Catadioptric cameras, Shape, normal and reflectance estimation

1 INTRODUCTION

T HE angular cone of light emitted by a surface often captures rich insights into the nature of its interaction with light. From a fairly mundane encoding of the surface's reflectance, be it Lambertian or specular, to more exotic phenomena like thinfilm interference and iridescence, measuring the radiance of light over a large solid angle—perhaps, even a hemisphere—can play a pivotal role in understanding shape and reflectance.

A light field camera, however, only measures a small portion of the cone of light emitted by a scene point. For example, a camera observing a scene point that is a meter away from a lens with a light gathering diameter of 5cm measures a small cone that has an angular spread of $\sim 3^{\circ}$. Even using a powerful microscope objective with a numerical aperture of, say, 0.5 only allows us to measure a cone of light with a spread of 60° . A multi-camera light stage, or perhaps a kaleidoscope, can be used to surround the object of interest to capture a larger cone of angles; yet, the number of measurements per scene point scales linearly in the number of cameras/viewpoints, which results in a sparse sampling of light rays over the angular cone of interest. Hence, current systems are woefully inadequate when it comes to providing a dense sampling of the wide angular cone of light emitted by a scene.

This paper proposes a novel design for acquisition of widebaseline light fields, i.e., a system that measures a large fraction of the cone of light emitted by scene points. We achieve this by using a standard light field camera that observes the scene of interest, not directly, but *through an ellipsoidal mirror*. This design relies on the observation that the ellipsoidal mirror morphs the wide cone of light from one of its foci to a much smaller cone at the other focus point; hence, placing the scene at one focus and the camera at the other allows us to acquire a wide-baseline light field. We show that the design is, in principle, capable of obtaining a larger cone of light than prior designs, often close to a hemisphere for many scene points.

The measurement of wide-baseline light fields can significantly enhance conventional uses of such signals while enabling a bevy of new capabilities. Like its narrow-baseline counterpart, wide-baseline light fields can be used to refocus the scene at different depths, reconstruct the shape of objects, and generate synthetic images from novel viewpoints. However, the larger angular cone can dramatically enhance these capabilities by providing extremely shallow depth of fields and the ability to image the backside of an opaque object. Additionally, wide-baseline light fields can provide truly novel capabilities. First, the wide angular cone allows us to capture high-frequency reflectance phenomena like iridescence that are not easily observable in other setups. Second, by using visibility as a physical cue, we can estimate the surface normals of scene points; such a capability, especially in the absence of any active illumination or strong assumptions on the scene illumination, is unique to our setup.

Contributions. This paper introduces a new methodology for capturing and analyzing wide-baseline light fields. In particular, we make the following key contributions.

- *Wide-baseline light field camera design*. We design a catadioptric device using a light field camera and an ellipsoidal mirror for capturing wide-baseline light fields.
- *Shape and normal estimation.* We describe a procedure for shape and normal estimation from wide-baseline light fields for reconstructing high-resolution 3D models while independently estimating the surface normal at each point.
- *Applications*. We explore applications traditional to light fields, namely, changing of the focus, perspective, as well as studying reflectance of scene points. The wider cones of light that our setup measures allows for some unique capabilities. For example, we can create focus stacks corresponding to planes with surface normal perpendicular to the optical axis of the setup. Reflectance captures with our setup can observe effects like iridescence, which are typically hard to measure.

We have released our code and data on the project website [1], to facilitate reproducibility and follow-up research.

The authors are with the Electrical and Computer Engineering Department at Carnegie Mellon University, Pittsburgh, PA, USA. E-mail: {mdezeeu2, saswin}@andrew.cmu.edu

Limitations. The proposed setup inherits a number of limitations, stemming from its use of mirrors and light field cameras. The ellipsoidal mirror presents a number of different constraints. Its size determines the size of objects that we can scan; this largely restricts our technique to small-sized objects-1 cm³ in our prototype. Further, the quality of the results also relies on the optical quality of the ellipsoidal mirror and the light field camera, both of which have limited commercial options. The use of ellipsoidal mirrors to relay the light field from one focus to another naturally results in a missing cone, centered around the axis joining the two foci; this missing cone results in non-traditional bokeh as well as reduced performance in normal estimation in some cases. A different set of limitations stem from the use of light-field cameras, which have implicit tradeoffs in achievable spatial and angular resolutions; we inherit these tradeoffs as well, and they affect the resolution and sample density compared to theoretical possibilities. Finally, the proposed shape and normal estimation algorithms use visibility analysis that relies on scene points having a supporting hyperplane; this analysis requires that the shapes under consideration are convex.

2 PRIOR WORK

Our work builds upon prior work devoted to spatio-angular sampling as well as the use of mirrors to enhance conventional lensbased imagers. We briefly discuss these topics.

2.1 Light Field Imaging

Microlens array-based light field cameras [2], [3] provide a compact optical design for resolving the scene focused on the image plane of a lens into spatial and angular dimensions. Light fields acquired from such devices have been studied extensively for refocusing [4] as well as estimation of shape and reflectance [5]; the interested reader is referred to [6] for a survey on light fields, and [7] for a study of space-angle tradeoffs with such designs.

Closely related to this paper is a body of work that aims to acquire wide *field of view* (FoV) light fields. Taguchi et al. [8] image an array of spherical mirrors with a conventional camera; here, each spherical mirror provides a wide FoV image with a slightly different perspective, thereby mimicking a camera array. Moving the camera axially in front of a symmetric mirror has also been shown [9] to produce wider FoV devices. Similar designs for expanding the FoV, using refractive optics, were considered in Dansereau et al. [10]. While these works are conceptually similar to ours, enhancing the FoV of the camera is significantly different from capturing a larger cone of light from scene points.

Light fields have also been studied in the context of microscopy [11], [12]. Unlike traditional photography, microscopes do gather a larger cone of light from scene points; crucially, the large *spatial magnification* of the microscope results in commensurate compression of the angular cone—a property that is shared with our design. However, there are two key differences. First, in a microscope, the cone of light measured from a scene point is largely centered around the optical axis of the device, which precludes the processing techniques that we introduce; in particular, estimating surface normals using visibility analysis is not possible with such a cone. Second, the cone of light that we measure is often larger and would require objective lenses that have impractically large numerical apertures.

Camera arrays [13] provide a different approach to acquire light fields; their large baseline provides a synthetic aperture that excels in focusing through complex occluders [14]. Closely related to camera arrays are light stages [15], used extensively for shape and reflectance acquisition, and kaleidoscopes, used for acquiring light fields [16] as well as 3D shape [17], [18], [19]. All of these techniques produce multiple view points—real or virtual—that fully surround an object; yet, each view point only measured light from a small cone and so we obtain a very sparse aliased sampling of the cone of light from any scene point.

2.2 Catadioptric Imaging

Our proposed approach falls under the class of catadioptric cameras, where mirrors are used to augment the refractive optics found in consumer devices. Baker and Nayar [20] study the family of mirror shapes that can alter the field of view of a perspective camera, and ellipsoidal mirrors are among the shapes that they analyze; our work can be interpreted as a specialized application of these ideas, when a scene and a light-field camera are placed in the two foci of such a mirror.

The large light collection area of mirrors has been used extensively to study reflectance properties of materials, using ellipsoidal [21], [22] as well as parabolic mirrors [23], [24], [25]. In particular, our use of ellipsoidal mirrors is motivated from Ghosh et al. [25] and Mukaigawa et al. [21], [22], where the BRDF of a material sample is measured over a large angle. However, there are notable differences that include the lack of active illumination in our setup, our use of a light field camera, as well as the computational techniques adopted by us. Further, in terms of scope, we look at extended 3D objects with the goal of estimating shape and normal recovery, as opposed to measurement of BRDF of a (typically) planar material sample.

3 OVERVIEW

We now describe our imaging setup and study properties of light fields in the context of an ellipsoidal mirror.

3.1 Imaging Setup

Our imaging setup uses a light field camera and an ellipsoidal mirror to image an object. Ellipsoids possess the property that all rays originating at one focal point and reflected off the surface of the ellipsoid will be reflected to the other focal point. By placing an object at the first focal point and the light field camera at the second, this provides a convenient starting point for mapping the light field at the object to the light field at the sensor. This arrangement, as we discuss next, allows us to capture a sufficiently broad angular cone from scene points.

The specific geometry of our imaging setup is largely driven by the market availability of image-quality ellipsoidal mirrors. Among the limited available options, we prioritized mirrors that could reflect to the widest possible cone of light measurable at the sensor to limit the size of the unmeasurable cone, with a second priority of having a large enough size to sample a larger spatial extent. The considerations led us to an ellipsoidal mirror with the geometry defined in Figure 1. As seen in the figure, an ellipsoidal mirror will reflect a wide cone of light at the object to narrow rays of light at the sensor. Since light rays that reflect off the upper region of the ellipsoid will be at wide angles on the sensor, they cannot be measured. So, our final imaging setup uses just a section of the ellipsoid instead. This results in the mapping shown on the right in Figure 1, where a wide-baseline light field that spans a



Fig. 1. Illustration of our imaging setup using an ellipsoidal mirror and a light field camera. The left hand side shows the setup with a full ellipsoidal mirror, which defines the overall geometry of the system. Our setup uses an ellipsoid with major and minor radii of a = 287.87mm and b = 135.47mm respectively. The two foci of the ellipsoid are spaced c = 508.0mm apart. Through the ellipsoid, wide angles at the object plane are mapped to narrow angles at the sensor plane while narrow angles at the object plane are mapped to wide angles. The right hand side shows the setup with a section of the ellipsoidal mirror that maps to the measurable cone of light for the light field camera. The corresponding light field mapping is shown in green, which illustrates the mapping of a wide-baseline light field (53.3° over the horizon $\times 180^{\circ}$ in azimuth) to a measurable light field (45.3° cone in diameter). The red cone in the center corresponds to the missing cone centered on the major axis of the ellipsoid.

 53.3° angle above the horizon and 180° in azimuth gets mapped to a 45.3° cone at the sensor, which is able to be measured using an existing light field camera.

The design of the microlens-based light field camera was driven by the need to measure the 45.3° cone of light propagating to the sensor based on the mirror geometry; this requires a microlens array with an *f*-number close to one. Based on this target and available optical elements, we choose an array with 2.2mm lenslet diameters with a 3mm focal length. To limit spatial blur, we mount a pinhole array aligned with the center of each lenslet on top of the microlens array; a detailed justification for the use of the pinhole array is presented in Section 3.3.

Using this imaging setup, we capture light fields like those shown in Figure 2. Under each lenset, we capture a different subset of the light field via a different view of the ellipsoidal mirror below. Different textural elements from the surface of the object are transformed into these different views. The figure shows these rendered light fields for a sphere and bunny model with different texture and reflectance characteristics. With this, we can analyze the angular coverage and the space-angle resolutions in more detail.

3.2 Angular Coverage

The use of the ellipsoidal mirror in our setup maps the widebaseline light field at the scene to a measurable light field captured by our sensor. There are some fundamental limitations to how much of the light field is transferred to a measurable space, which determines the coverage of our imaging setup.

We start with an analysis of a 2D version of our setup. Defining the two foci of an ellipse at $y_1 = -254$ mm and $y_2 = 254$ mm, we can use simple ray tracing to propagate an object-plane light ray (x_1, θ_1) to its corresponding sensor-plane light ray (x_2, θ_2) .



Fig. 2. Example images rendered from a simulation of our imaging setup, including a close-up of two sub-aperture views for each scene. The black disk at the center of each sub-aperture is a consequence of the missing cone, highlighted in Figure 1.

$$x_{2} = f_{1}(x_{1}, \theta_{1})$$

$$= \frac{y_{2} - y_{1} - \hat{t}(x_{1}, \theta_{1})\sin\theta_{1}}{\tan\theta_{2}} + x_{1} + \hat{t}(x_{1}, \theta_{1})\cos\theta_{1} \quad (1)$$

$$\theta_{2} = f_{2}(x_{1}, \theta_{1})$$

$$= \pi - \theta_1 + 2 \tan^{-1} \left(\frac{b^2 \left(y_1 + \hat{t}(x_1, \theta_1) \sin \theta_1 \right)}{a^2 \left(x_1 + \hat{t}(x_1, \theta_1) \cos \theta_1 \right)} \right), \quad (2)$$

with \hat{t} and supporting functions defined as

$$\widehat{t}(x_1,\theta_1) = \frac{\widehat{\gamma}(x_1,\theta_1) + ab\sqrt{\widehat{\alpha}(\theta_1) + \widehat{\beta}(x_1,\theta_1)}}{\widehat{\alpha}(\theta_1)}, \qquad (3)$$

$$\widehat{\alpha}(\theta_1) = a^2 \cos^2 \theta_1 + b^2 \sin^2 \theta_1, \tag{4}$$

$$\widehat{\beta}(x_1, \theta_1) = 2x_1 y_1 \cos \theta_1 - x_1^2 \sin^2 \theta_1 - y_1^2 \cos^2 \theta_1, \quad (5)$$

$$\widehat{\gamma}(x_1,\theta_1) = -a^2 x_1 \cos \theta_1 - b^2 y_1 \sin \theta_1.$$
(6)

Here, a and b are the major and minor radii of the ellipse, respectively. Please refer to the supplemental material for a detailed derivation. Figure 3 shows a visualization of this propagation for both a full ellipse and a section of the ellipse, which corresponds to the section of the ellipsoid used in our imaging setup. The objectplane light field is color coded, with brightness corresponding to spatial location and color corresponding to angle. The plot shows how those color-coded light rays propagate to the sensor plane. The measurable region, assuming a 1 inch sensor and a microlens array capable of imaging a 45° cone of light, is shown on the plot. This highlights the key benefit of the ellipsoidal mirror, which maps the grazing angles (close to 0° [red] and 180° [blue]) to the measurable region at the expense of other angles.

Comparison to common light field setups. It is illustrative to compare the sampling of our setup with those achieved by using one or multiple conventional light field cameras. Figure 4 visualizes a setup where two light field cameras are placed 12.5cm from the sample with a light collection area of 5cm; these numbers approximately correspond to the shortest focusing distance and



Fig. 3. Plot of the light field propagation from the object plane to the measurement plane for a complete (lower left) and partial (lower right) ellipse. Rays at the object plane are color coded according the the plot in the top left, and the corresponding mapping to the measurement plane is the location of that color on the lower plots. The measurable region, assuming a 1" sensor and a 45° measurable cone, is marked on the axes and highlighted. The actual section of the ellipse used in our imaging system omits primarily the unmeasured light rays, resulting in the mapping in the bottom right. This section preserves the wide angular range that we are looking for at the expense of the central cone of light most commonly captured in other setups.



Fig. 4. Comparison of the light field sampling of our setup versus that of a light field camera pair. The left shows a plot of the coverage of the hemisphere of outgoing rays for each configuration. The right shows the light field radiance plots of this hemisphere for each set of samples.

entry aperture of the Lytro Illum. We can immediately observe the difference in angular coverage between the two setups; for example, the proposed imaging setup captures light rays along the horizon that the light field cameras miss. This difference in coverage affects subsequent processing of the light field for shape recovery as well as view point synthesis; specifically, the broader sampling in our setup enables the estimation of surface normals (Section 4.1.1) as well as viewpoint synthesis along orthogonal planes (Figure 20).

The missing central cone. A drawback of using an ellipsoidal mirror to remap wide-baseline light fields is that the span of angles that we measure is not continuous. Specifically, the imaging setup does not measure the cone of light from the scene that is centered around the major axis of the ellipsoid. This missing cone, marked in Figure 1, results in the dark central disk in the sub-aperture images in Figure 2. The missing central cone affects the performance of the shape estimation and rendering algorithms that we describe later. Perhaps ironically, this is the very cone of light that is (easily) captured by a traditional light field camera via the use of its objective lens.



Fig. 5. Plots of the object-plane angular and spatial resolution at $x_1 = 0$ mm as a function of θ_1 . The resolution is strongly dependent on the spatial blur resulting from the chosen aperture size Δx_2 , where different sizes are represented in the different columns. The top row plots the absolute value of the object-plane spatial resolution $|\Delta x_1|$ while the bottom row plots the absolute value of the object-plane angular resolution $|\Delta \theta_1|$. In our formulation, the smaller the resolution the better.

3.3 Analysis of Spatial, Angular, and Depth Resolution

The limitations of resolution on the sensor propagate to the measured light field in particular ways. We perform gradient analysis on the propagated light field formulation defined in the previous section to quantify how changes at the sensor plane (corresponding to the achievable spatial and angular resolution) propagate back to the object plane. Based on the expression in (1, 2), we derive the changes in x_2 and θ_2 based on changes in x_1 and θ_1 as:

$$\begin{bmatrix} \frac{\Delta x_1}{2} \\ \frac{\Delta \theta_1}{2} \end{bmatrix} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial \theta_1} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial \theta_1} \end{bmatrix}^{-1} \begin{bmatrix} \Delta x_2 \\ \Delta \theta_2 \end{bmatrix}$$
(7)

To evaluate this expression, we define Δx_2 as the size of the aperture used over the microlens array, for this indicates the size of one spatial "pixel". We test this value over multiple possible aperture sizes. We approximate $\Delta \theta_2$ as the pixel pitch of the sensor divided by the focal length of the microlens: $\Delta \theta_2 \approx \frac{5.94 \mu m}{3 mm} = 0.00198 = 0.113^\circ$. Using these values and computing the partial derivatives of f_1 and f_2 based on their definitions in (1, 2), we can analyze the resulting object-plane resolution for various portions of our object-plane light field.

Figure 5 plots different object-plane resolutions for three aperture sizes defined in the figure. Our imaging setup only captures outgoing angles up to about 50° , which corresponds to approximately $0^{\circ} - 50^{\circ}$ and $130^{\circ} - 180^{\circ}$ on the plots; this does not include angles close to 90° that observe worse spatial resolution. Focusing on these regions of interest, we can see that using no aperture on the microlens array results in very poor resolution in both space and angle. Adding the $200\mu m$ aperture improves the angular resolution to between 0.1° and 0.5° and the spatial resolution to less than $200 \mu m$. Adding the $10 \mu m$ aperture maintains a similar angular resolution and improves the spatial resolution further to less than $10\mu m$, but this comes with the risk of significant diffraction blur that is not modeled by ray tracing. Figure 6 shows a simulated image capture of a diffuse textured sphere with no aperture and with a $200\mu m$ aperture. The added aperture prevents the blurring of textural elements in the image capture that are needed for various processing techniques.

We also analyze the theoretical benefits to depth resolution provided by wide-baseline light fields. Figure 7 illustrates the



Fig. 6. Simulated image captures with no aperture (right) and with a $200\mu m$ aperture (left). Adding the aperture preserves textural details in the image capture, as evident in the magnified areas shown.



Fig. 7. Illustration of the derivation of depth resolution as a function of measured angular range. We test a point (in red) to see if it belongs on the surface along the bottom when it is placed at distance *d* above the surface. If we are able to measure a light field with angular range θ , then we can produce a virtual sensor plane of size *s* at distance *f* from the point. With this geometry, a texture element of size *T* on the surface will project to size b_T on the virtual sensor.

relevant terms. The goal is to determine whether a given point, shown in red, belongs on the surface of a textured object, shown as the plane along the bottom. The point is set at distance d from the surface. We define s as the size of a virtual sensor positioned f away from the point, where s is determined by the angular range θ of the measured light field. The size of a surface texture element T is mapped to an area on the sensor of size b_T .

We can see that as d decreases with all the other variables fixed, the projected texture size b_T would increase. We determine that a point belongs on the surface when all of the samples on our virtual sensor correspond to the same point, which occurs when $b_T = s$. At this stage, the angle subtended by a texture element of size T on the test point d units away is equal to θ , the angular span of the light field. We can now find d_{max} , the maximum error in our depth estimation, as:

$$d_{max} = \frac{T}{2\tan\left(\frac{\theta}{2}\right)} \tag{8}$$

As θ approaches a full hemisphere of 180° , d_{max} approaches 0. So, if we are able to measure a full hemisphere of light above every surface point, there are no theoretical limits to the achievable depth resolution with our method. Instead, the achievable resolution is limited only by the span of the angular cone that we measure.

4 3D POINT AND NORMAL ESTIMATION

We now design algorithms for recovering shape, in the form of 3D points and associated surface normals.

4.1 Shape Reconstruction

We approach shape reconstruction from the mindset of scanning the volume, one voxel at a time, and establishing whether or not a voxel lies at (or near) the surface of an object. Our approach for determining this relies on an observation on determining the tangent plane to the scene point from light field measurements.

Observation. Consider the example in Figure 8 where we visualize the radiance from a point that is on the surface of a textured sphere; this point under consideration has an amber-color diffuse albedo and is mildly specular. The figure shows the radiance over a full sphere of outgoing directions, i.e., in every possible outgoing direction (θ, ϕ) , where θ and ϕ are azimuth and elevation angles about the vertical line. This radiance image has two clear segments: a near constant region that corresponds to light that is emitted by the surface, and a textured region that corresponds to light from other scene points that occlude the point under consideration when observed along the corresponding direction. The boundary between the two corresponds to the directions on the tangent plane to the surface at the point under consideration, and maps to a circle on the unit sphere of directions. When we consider a point that is even slightly off the surface-either above or below—as seen in Figure 8 (bottom), this behavior changes completely; in particular, the radiance map is not crisply split into two regions with differing behavior, nor is the separating boundary between them guaranteed to be a circle. The observations do require the sufficiently rich textures and fairly mild assumptions on the specularity of the surface. More importantly, this observation provides the basis for identifying between voxels that contain valid scene points and those that do not; intriguingly, they also allow us to estimate the surface normals.

In essence, our approach for 3D point and normal estimation relies on whether or not we can find a valid surface normal at each candidate voxel; failure to find such a normal indicates that there is no valid 3D point in the voxel.

4.1.1 Normal Estimation

To solve for the surface normal \vec{n} at each point, we utilize the visibility clues that we gain from the wide-baseline light field. In general, an outgoing direction \vec{v} from a point will belong to the hemisphere around that point's surface normal if $\vec{v}^{T}\vec{n} > 0$. For convex objects, this hemisphere is the set of direction from which the point is visible, which means that the radiance of sampled rays from within this hemisphere will be samples of the reflectance at that point. This hemisphere is visualized by the constant region in the upper half of the plot in Figure 8.

Conversely, directions \vec{v} such that $\vec{v}^{\top}\vec{n} < 0$ in the lower region correspond to sampled directions at which the surface point is not visible. These directions instead sample the radiance of various points on the other side of the object.

If the object is textured, as it is in this case, we can describe these two regions as "low variance" (for the region above the visibility horizon plane) and "high variance" (for the region below the visibility horizon plane). The samples in the low variance region all correspond to the reflectance at a single point, which has a large amount of consistency across outgoing angles. The samples in the high variance region sample the reflectance at a variety of surface points, which all could have unique reflectance properties. Therefore, we can consider the plane $\vec{v}^{T}\vec{n} = 0$ as corresponding to the visibility horizon. In Figure 8, the visibility horizon plane corresponds to the transition from the low variance



Fig. 8. Plot of the entire light field of radiance measured along each outgoing direction (θ, ϕ) (defined in the upper right) from a single point either on (top), above (bottom left), or below (bottom right) the textured sphere shown in the top right. The ray diagram in the top left shows how different directional samples from this single point map to different radiance measurements. When the point is on the surface, the set of all sampled light rays over the hemisphere above that point will correspond to the reflectance of that point. This results in the consistent reflectance observed in the upper half of the topmost plot. This scene was illuminated using four directional light sources, which appear as the four specular highlights in the plot. Sampled light rays from the hemisphere below the surface point will actually be samples from various points on the other side of the sphere, which results in the varying, checkered reflectance pattern observed in the lower half of the plot. The transition between these two regions corresponds to the visibility horizon. When not on the surface, this transition is not defined and so we cannot estimate a surface normal.



Fig. 9. Plot of the portion of the light field measured by our imaging setup as radiance measured along each outgoing direction (θ, ϕ) from a single point, similar to Figure 8. The visibility horizon plane, plotted in green, still adheres to the boundary between high and low variance regions.

region in the upper part of the plot to the high variance region in the lower part of the plot.

With our imaging setup, we are not actually able to sample every outgoing ray from every point in the scene. Instead, for a given point in the scene, the outgoing rays that we can measure would produce a plot like the one in Figure 9. The visibility horizon plane, plotted in green, is still observable as the boundary between high and low variance regions. To detect this horizon plane for every point analyzed in the scene, we follow a four-step algorithm, with each step visualized with an example in Figure 10.

- First, we color-normalize our measurements, i.e., to ensure that effects of shading are not considered as high variance effects, we normalize every measurement by the sum of its color channels.
- Second, we compute a local variance metric for each sampled direction. This metric takes samples from the neighborhood around a given direction and computes the sum of the variances of the color-normalized red, green, and blue channels.



Fig. 10. Visualization of the four stages of the normal estimation pipeline. Each plot is consistent with the (θ, ϕ) axes defined previously. A: color normalized reflectance plot. B: local variance metric computed for each direction (brighter corresponds to higher variance). C: variance labels (green is low variance, red is high variance). D: visualization of the fitted visibility horizon plane (red) along with ground truth visibility horizon plane (green) over the original reflectance plot.



Fig. 11. Network trained to classify high and low variance regions.

- Third, we assign a label y for each direction based on a threshold chosen empirically. We set y = 1 for low variance points and y = 0 for high variance points.
- Fourth, we train a simple classifier to separate the low and high variance regions. The classifier network structure is shown in Figure 11. The output of the network is a class probability predictor *p*, defined as

$$p = \frac{1}{1 + e^{-\vec{v}^{\top}\vec{n}}}.$$
 (9)

The connectivity of the network ensures that the input to the logistic function is $\vec{v}^{\top}\vec{n}$, so the function will output a class probability based on whether that dot product is greater than or less than zero, which corresponds to the visibility horizon plane classification discussed previously. The class probability $p \in [0, 1]$ will be closer to 0 when $\vec{v}^{\top}\vec{n} < 0$ and closer to 1 when $\vec{v}^{\top}\vec{n} > 0$. We train our network using a standard binary cross-entropy loss function. The weights of the network are the elements of the surface normal estimate \vec{n} once normalized to a unit vector.

This algorithm will estimate a surface normal for any point in the scene. For points on the surface of an object, the visibility horizon plane corresponds to the delineation of high and low variance regions in the reflectance. For points that do not fall on the surface, such a plane does not exist. Therefore, the ability to detect and estimate an effective visibility horizon plane provides depth cues that can be used to generate a point cloud of the object.

4.1.2 Depth Reconstruction

We use the surface normal estimation procedure, described earlier, to determine whether or not a point belongs to the surface. Given a 3D point under consideration, we first identify visible light rays



Fig. 12. Histogram of the total variance of visible rays as a function of distance from the surface. Moving away from the surface in either direction results in a significant increase in variance, so we set a variance threshold to detect point that are on or very close to the surface.

from it, using the variance metric. Using just these visible rays, we can compute the total variance of rays above the visibility horizon. If the total variance is below a certain threshold, then we denote that point as belonging on the surface. If the overall variance is above that threshold, we reject the point as not belonging to the surface. Similarly, if the normal estimation algorithm does not find a large enough set of rays that are considered visible, we reject the point as likely falling within the object. In either case, these tests determine whether or not the normal estimation algorithm produced a visibility horizon plane that successful delineates a significant region of low variance reflectance measurements. The variance of the visible region is ultimately a function of a point's distance from the surface. Figure 12 shows a histogram of the variance of visible rays for a textured sphere as a function of the distance from the surface. We see that the relationship between distance from the surface and variance holds well, for moving a small distance away from the surface produces a large jump in variance relative to points on or very close to the surface.

5 SIMULATED RESULTS

We evaluate the proposed setup and associated algorithms with simulated measurements.

5.1 Shape Reconstruction

We perform shape estimation using the procedure outlined above. The output of the algorithm is a point cloud where each point has an associated surface normal and computed variance. To improve the sampling of the light field, we capture results from a 6×6 grid of shifted camera positions. This aids the variance computations by providing more data, which helps improve the reconstructions. However, capturing with a single exposure can still produce quality results. Figure 13 shows a side-by-side comparison of point clouds generated at 150μ m resolution for a single exposure and for the multiple exposure method. The multiple exposures reduce some artifacts and produce slightly better results, but the single exposure still captures the shape qualities of the object.

In both cases, we observe that gaps in the point clouds occur due to missed points in the first phase of the coarse-to-fine volumetric sweep that we perform. These gaps most often occur in areas where the convexity of the object least holds, for that



Fig. 13. Comparison of point clouds produced at $150\mu m$ resolution for data from a single exposure and a 6×6 grid of exposures. The color map is based on the depth of the points.

reduces the size of the cone of light that would be considered low variance for points on the surface. Occasionally, these gaps occur in other areas as well, which occurs due to the coarse sampled point that is nearest the surface in that area being very near to a textural transition, which can introduce more variance.

We run our algorithm on four objects: a sphere, bunny, armadillo, and dragon. Due to the wide angle of measured light, we are able to observe points on the far side of these objects from the camera. However, depending on the shape of the object, the amount of sampling on the backside may be limited. In particular, the bunny model has a large area on the backside with a surface orientation directly away from the camera, so we are unable to measure sufficient samples there. For the sake of visualization, we merge point clouds generate from setups configured above and below the object to fill in the full reconstruction, with each view producing a high quality reconstruction of the side of the object facing the camera. The sphere, armadillo, and dragon models are sufficiently sampled on the backside to fill out a 3D reconstruction, so these objects are reconstructed from the single view.

Figure 14 shows reconstruction results for these four objects captured in simulation. The objects were rendered with high frequency diffuse textures and analyzed down to a point cloud resolution of $50 \mu m$. The top three rows of the figure show the ground truth normal map visualized on the objects followed by two views of the estimated normals visualized on the reconstructed surfaces. We observe that while some of the high-frequency information from the ground truth surface normals is not preserved, the overall surface orientations are effectively estimated.

The bottom three rows of the figure show ground truth objects from one view and reconstructed objects from two different views. The reconstructed objects were generated from the point cloud locations and normals using MeshLab. We observe that for relatively simple, mostly convex shapes like the sphere and the bunny, our reconstructions are quite accurate. The method begins to break down as we try to resolve the fine shape details and nonconvex areas of the armadillo and dragon, but the overall shape characteristics are reconstructed still.

5.2 Refocusing

The wide baseline of measured light enables us to generate focal stacks along different directions. Figure 15 shows three sweep directions for a two-color diffuse textured sphere. Here, our setup



Fig. 14. Shape reconstructions for the sphere, bunny, armadillo, and dragon models. Using the point clouds with per-point surface normals, we use MeshLab to generate reconstructed surfaces. To visualize the original surface normal estimates, we reassign the surface normals for each vertex in the reconstructed mesh to equal a weighted sum of the normal estimates for the ten nearest points in the original point cloud. The top row visualizes the ground truth surface normals. The second and third rows show two views of each reconstructed object with the estimated surface normals visualized on the surface. The fourth row shows ground truth renderings of the different models. The final two rows show two views of rendered images for the reconstructed meshes. We also compute the average reconstruction error of the original estimated point cloud for each of the objects: sphere - 0.116mm; bunny - 0.584mm; armadillo - 0.125mm; dragon - 0.093mm. Each measured object fits snugly within a 1 cm^3 volume, which means our error is approximately in the range of 1-5%.



Fig. 15. Simulated refocusing of the sphere model with a diffuse texture at different depths. Each row corresponds to a different focal sweep direction as visualized by the graphics on the left. All directions are with respect to the light field camera direction, which is labeled. The hourglass artifact in the last row is a consequence of the central subaperture view being aligned with the ellipsoid focal point; in this case, all rays from this subaperture map exactly to the other focal point, while neighboring sub-apertures map to points some distance away.

allows for views around the backside of the object from the camera. In the sweep from the side, we get better sampling in the region toward the camera (on the top of the images), but we still focus on the texture at the bottom of the sphere.

5.3 Viewpoint Synthesis

In addition to refocusing, we can resample our light field to generate synthetic views of our objects. However, the sparse spatial sampling of our light field in a single exposure hampers our ability to produce high quality pinhole views. Similar to shape reconstruction though, we can capture a set of data with shifted camera positions to fill in this sparse sampling. Figure 16 shows the results of four different views generated for different objects with different reflectance properties. Note in particular the captured variations of the specular highlights of the different objects. The mixed textures are a random assortment of diffuse and specular materials based on those captured from the MERL dataset [26]. In addition, the bottom row shows a sphere modeled with an iridescent thin film coating. The different synthesized viewpoints observe a wide selection of specular angles that cause the specular highlight to shift in color according to the model.

6 REAL RESULTS

We now discuss our lab prototype and explore shape estimation as well as viewpoint synthesis results on real objects.

6.1 Lab Prototype

To implement this imaging setup, we use components based on the outline of the imaging setup defined earlier. For the light field camera, we use a Nikon Z6 camera with a hexagonally-packed microlens array mounted one focal length away from the sensor. Each lenslet has a diameter of 2.2mm and a focal length of 3mm. We additionally mount an aperture array above the microlens array, where each 200μ m pinhole of the aperture aligns with the center of each lenslet. This camera as a whole is mounted above the object at the approximate location of the second focal point of the ellipsoidal mirror. The section of ellipsoidal mirror is mounted



Fig. 16. Simulated viewpoint synthesis of the sphere and bunny models with different textures and reflectances. The gold and mixed texture bunnies were illuminated with four directional light sources. The sphere model was illuminated with a single light source. The iridescent sphere was modeled using a thin film reflectance model, which produces angle-dependent changes in reflected wavelengths in the specular highlights.



Fig. 17. Images of the real imaging setup. On the left is the entire camera, with the ellipsoidal mirror section below, the light field camera above in the center, and a helper camera and projector mounted around the apparatus for calibration. On the right are two views from the camera toward the mirror for two different objects. Each object is set on top of a vertical translation stage at the focal point. A list of components is provided in the supplemental material.

around the object, which is place at the approximate location of the first focal point. The full setup is shown in Figure 17. For illumination, we create a small spotlight to illuminate the object without adding specular artifacts by reflecting off the mirror. We do not calibrate or control the location of this light source. The figure also shows an additional camera and projector added to the setup to aid in calibration, which will discuss next.

Calibration. In order to know the correspondence between measured pixels on our camera and light rays at the object plane, we need to calibrate our imaging setup. However, calibration of this setup is significantly complicated due to imperfections in the shape of the ellipsoidal mirror, which was designed primarily to concentrate light from one focus to another and not as a precision imaging device.

To achieve our calibration, we add a helper camera focused



Fig. 18. Three views of the reconstructed point cloud of a textured plane (a bayer pattern printed on paper) with computed surface normals. These images show how the blur in the system allows proper sampling mostly just in the center of each texture element, as the sampling gaps follow a regular grid. The estimated surface normals are all fairly consistent with the proper orientation of the plane, but they do tend to be skewed toward whichever end of the volume the point is on. In the bottom right, we visualize the point cloud along with a fitted plane.

at the object plane along with a laser projector aimed at the object plane. The helper camera is calibrated using standard camera calibration techniques. We place a spectralon mounted on a vertical translation stage at the object plane. The first step of the calibration pipeline is to place a checkerboard on the spectralon and use the helper camera to get the plane's position in 3D; we project a sparse set of points and estimate its homography to the helper camera. Next, we project a series of row and column graycode patterns onto the spectralon and capture images from the light field camera for each. The codes provide a mapping from light field camera pixel to projector pixel, which in turn can be mapped to the 3D plane via the helper camera.

This pipeline is then repeated for a series of depths by adjusting the translation stage. At this point, for each pixel on the light field camera we have a set of 3D points corresponding to the different depth planes. We fit a ray through these points using RANSAC to achieve the target of pixel-light ray correspondence.

6.2 Real Results

With the lab prototype, we can capture single-exposure images of wide-baseline light fields. We use this data to perform the same set of processing that we showed in simulation.

Shape Reconstruction. The blur that persists in our real imaging system prevents the processing of more complex objects. However, Figure 18 shows a proof of concept of this processing technique on a planar textured surface. Using the joint depth and normal estimation, we produce the point cloud shown in the figure.

Refocusing. We generate focal stacks with a small depth of field in the same way that we did for the simulated results. Figure 19 shows results for two objects, all generated with a vertical focal sweep from under the object surface up toward the camera. The first object is a ThorLabs resolution chart printed on glass. The first two rows of the figure show data captured in two different areas of the chart. For both locations, we observe two apparent focal planes: first in column (b), then in column (d). While column (b) seems less in focus than column (d), we note that it is more in focus than the focal plane between the two in column



Fig. 19. Focal sweeps of real data captured of a ThorLabs resolution chart and a blue die. Depths are swept from lower to higher in the direction of the light field camera.



Fig. 20. Focal sweeps in vertical (top) and horizontal (bottom) directions, as illustrated on the left. We are able to resolve different faces on the die through the different directions.

(c). This dual focusing is caused by the shadow cast from the printed pattern through the glass to the plane on which the glass is resting. Column (b) is the depth focused on this shadow, which is inherently less sharp than the original pattern and is shifted slightly in one direction according to the direction of the light source. Column (d) is focused on the pattern printed on the glass, which is sharpest. On either side of these focused planes, the focal sweep quickly defocuses. The final object shown is a blue die. We see that we come clearly into focus in column (d), but fall into defocus blur in either direction from that point.

As we showed in simulation, we can also perform focal sweeps in different directions as well. Figure 20 shows two focal sweep directions for a stack of two small dice, particularly focused on the top die where we have more samples. Sweeping from inside the die up through the top toward the camera, we come into focus in the fourth frame. The second frame is focused on two of the dots on the side of the die, so the in focus area resolves to a small strip from this view. The sweep from the side of the die brings a different face of the cube into focus and effectively illustrates the narrow depth of field. Since we are not sweeping perfectly square to the side of the cube, one part of the face comes into focus before the other, as shown in the third and fourth frames. This is true for the first two frames as well, where different parts of the row of three dots on top come into focus in each frame.

Viewpoint Synthesis. In our simulations, we improved the quality of viewpoint synthesis results by capturing a set of images with a shifted camera. This produced high-quality pinhole images over a wide range of viewpoints. With our physical setup, we capture just a single image with a stationary camera. This does not provide the dense sampling needed for these sharp synthetic views. Instead, we simulate a lens-based synthetic camera and apply interpolation to the rays we have sampled to fill in the missing gaps. This produces



Fig. 21. Viewpoint synthesis for a blue die (top row), a small green die (middle row), and a blue resistor (bottom row). Each view reflects the angular changes in reflected light based on a single illuminant.



Fig. 22. Real data captures of iridescent objects (iridescent rock, top, and compact disc, bottom). Different sampled light rays have different colors based on the angle at which the object is observed.

blurrier images than what we showed in simulation, but we are able to produce these results from a single exposure. Figure 21 shows results for four different objects. In particular, note how the different views capture the different angular responses to the single light source. The specular highlight in the third column for blue die and the different illuminated or shadowed side faces on the green illustrate this behavior effectively.

Iridescence. Beyond the developments and results of the prior processing techniques, we can show that our wide-baseline light field camera can capture complex reflectance phenomenon such as iridescence. Figure 22 shows the data captured with our setup of two different iridescent objects: an iridescent rock and a compact disc (CD). The close-up of a single aperture view reveals the effective sampling of iridescent qualities, as the rock shows reflections of green, blue, and purple in different directions and the CD shows specular highlights in a broad range of colors. The spectrum of the observed reflectance for each is dependent on the angle of observation.

We can visualize iridescent reflectance through the reflected light in every direction from a point on the surface. Since we do not have a ground truth shape and 3D location of these objects and the complex reflectance prevents us from computing depth from variance, we estimate that the object is approximately planar in the area considered and passes through the focal point of the ellipsoid. We can therefore generate data at points around the focal point on this plane. We plot one of these points for each object to visualize in Figure 23. Similar to before, we plot the measured radiance for outgoing angles θ and ϕ from the point under consideration. Here, we limit the range of ϕ to the observable range on the ellipsoid, and θ covers all angles from 0 to 2π . The figure reveals that as



Fig. 23. Reflectance plots showing the light sampled in different outgoing directions (θ, ϕ) for a point on the surface of an iridescent rock and a compact disc. The iridescent nature of these objects is effectively captured over a broad range of angles.

we observe these objects from different directions, the spectrum of the reflected light changes; the results shown are speckled due to the sparser sampling introduced by the pinhole array. For the iridescent rock, we observe reflectance shifting from green to purple for different angles of θ . For the CD, we note many shifts in color over small changes of θ and ϕ . These iridescent reflectance effects reveal themselves over the broad baseline of samples that our proposed setup captures.

7 CONCLUSION AND DISCUSSIONS

We introduced a novel imaging device that senses wide-baseline light fields through the mapping properties of an ellipsoidal mirror. We showed how this data enhances the capabilities of standard light field processing techniques by achieving extremely small depth of fields and increasing the area of the object that we can observe. Wide-baseline light fields also open the door to a new set of capabilities. By using visibility cues, we detailed a procedure for estimating surface normals of scene points and integrated this ability into a shape reconstruction algorithm. We also demonstrated the ability to capture high-frequency reflectance phenomena like iridescence. In total, our proposed camera design pushes into new territory for light field acquisition and analysis.

Assumptions on object shape. A key highlight of wide-baseline light fields is their ability to resolve surface normals from a visibility analysis. This relies on the existence of a supporting hyperplane at a scene point, such that there are no self-occlusions in the cone of light emitted on one side of the hyperplane. This property is satisfied by a convex object. In practice, while this technique does work on non-convex shapes as seen in Figure 14, the reconstruction is worse in regions that are locally non-convex.

Effect of the missing cone of light. As highlighted earlier, light field capture using ellipsoidal mirrors is incapable of measuring the cone of light centered around the major axis of the ellipsoid. This has important consequences both in shape estimation as well as image-based rendering.

Our shape estimation techniques rely on the large angular span of the measurements, as opposed to completeness of the sampling in that range. This large angular span allows us to obtain samples on either side of the visibility horizon plane—which is critical for accurate determination of the surface normal as well as depth. This is visualized in Figure 9, which shows successful normal recovery from observing a small portion of the light directly on either side of the visibility horizon plane.



Fig. 24. Focus stacks for a plane near the focal point with five green dots (gamma-corrected). The defocus of these dots reveals the unconventional bokeh resulting from our setup and sampling. This varies depending on the orientation relative to the mirror, as each row designates.

The missing cone also results in some artifacts (seen in Figures 20 and 21) as the rendered images have unusual bokeh, which we simulate in Figure 24 for focus stacks with different orientations. The missing cone also restricts the view points where we can render perspective images. However, this is still less restrictive than what's possible with a standard light field camera, where the viewpoints are limited to the very narrow measured cone.

Increasing angular coverage. It is possible to incorporate additional optics to capture the missing cone of light. For example, we could use beamsplitters to optically collocate a second light field camera with a macro lens that captures *only* the central cone. Coverage can also be increased with the use of a larger ellipsoidal mirror, and making multiple captures, perhaps using the wide FOV device proposed in [10]. We can double the angular coverage by mirroring our current setup, placing a second ellipsoidal mirror that shares the same optical axis and focus at the object but captures the cone of light emitted below the focus. As is to be expected, all of these approaches pose engineering challenges in mounting and calibration.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under the grants 1801382 and 1652569. We also thank Prof. Kyros Kutulakos for early discussions that inspired this work, and Mr. Byeongjoo Ahn for help with quantitative evaluation.

REFERENCES

- M. De Zeeuw and A. C. Sankaranarayanan. (2022) Project page: Wide-baseline light fields using ellipsoidal mirrors. [Online]. Available: https://github.com/Image-Science-Lab-cmu/WideBaselineLF_ICCP22
- [2] G. Lippmann, "Epreuves reversibles donnant la sensation du relief," J. Phys. Theor. Appl., vol. 7, no. 1, pp. 821–825, 1908.
- [3] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, "Light field photography with a hand-held plenoptic camera," Stanford university, Research Report CSTR 2005-02, 2005.
- [4] R. Ng, "Fourier slice photography," ACM Trans. Graphics, vol. 24, no. 3, p. 735–744, jul 2005.
- [5] T.-C. Wang, M. Chandraker, A. A. Efros, and R. Ramamoorthi, "Svbrdfinvariant shape and reflectance estimation from light-field cameras," in *CVPR*, 2016.
- [6] M. Levoy, "Light fields and computational imaging," *IEEE Computer*, vol. 39, no. 8, pp. 46–55, 2006.
- [7] T. G. Georgiev, K. C. Zheng, B. Curless, D. Salesin, S. K. Nayar, and C. Intwala, "Spatio-angular resolution tradeoffs in integral photography." *Rendering Techniques*, vol. 2006, no. 263-272, p. 21, 2006.
- [8] Y. Taguchi, A. Agrawal, A. Veeraraghavan, S. Ramalingam, and R. Raskar, "Axial-cones: Modeling spherical catadioptric cameras for wide-angle light field rendering," *ACM Trans. Graphics*, vol. 29, no. 6, p. 172, 2010.
- [9] Y. Taguchi, A. Agrawal, S. Ramalingam, and A. Veeraraghavan, "Axial light field for curved mirrors: Reflect your perspective, widen your view," in *CVPR*, 2010.

- [10] D. G. Dansereau, G. Schuster, J. Ford, and G. Wetzstein, "A wide-fieldof-view monocentric light field camera," in CVPR, 2017.
- [11] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz, "Light field microscopy," ACM Trans. Graphics, vol. 25, no. 3, p. 924–934, jul 2006.
- [12] X. Lin, J. Wu, G. Zheng, and Q. Dai, "Camera array based light field microscopy," *Bio. Op. Exp.*, vol. 6, no. 9, pp. 3179–3189, 2015.
- [13] B. Wilburn, N. Joshi, V. Vaish, E.-V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy, "High performance imaging using large camera arrays," *ACM Trans. Graphics*, vol. 24, no. 3, p. 765–776, jul 2005.
- [14] V. Vaish, M. Levoy, R. Szeliski, C. L. Zitnick, and S. B. Kang, "Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures," in *CVPR*, 2006.
- [15] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *SIG-GRAPH*, 2000, p. 145–156.
- [16] A. Manakov, J. Restrepo, O. Klehm, R. Hegedus, E. Eisemann, H.-P. Seidel, and I. Ihrke, "A reconfigurable camera add-on for high dynamic range, multispectral, polarization, and light-field imaging," ACM Trans. Graphics, vol. 32, no. 4, pp. 47–1, 2013.
- [17] R. Xu, M. Gupta, and S. K. Nayar, "Trapping light for time of flight," in *CVPR*, 2018.
- [18] D. Lanman, D. Crispell, and G. Taubin, "Surround structured lighting: 3d scanning with orthographic illumination," *Computer Vision and Image Understanding*, vol. 113, no. 11, pp. 1107–1117, 2009.
- [19] B. Ahn, I. Gkioulekas, and A. C. Sankaranarayanan, "Kaleidoscopic structured light," ACM Trans. Graphics, vol. 40, no. 6, pp. 1–15, 2021.
- [20] S. Baker and S. K. Nayar, "A theory of catadioptric image formation," in *ICCV*, 1998.
- [21] Y. Mukaigawa, K. Sumino, and Y. Yagi, "Rapid BRDF measurement using an ellipsoidal mirror and a projector," *IPSJ Transactions on Computer Vision and Applications*, vol. 1, pp. 21–32, 2009.
- [22] —, "High-speed measurement of brdf using an ellipsoidal mirror and a projector," in CVPR, 2007.
- [23] H. Zhang, K. Dana, and K. Nishino, "Reflectance hashing for material recognition," in CVPR, 2015.
- [24] K. J. Dana and J. Wang, "Device for convenient measurement of spatially varying bidirectional reflectance," JOSA A, vol. 21, no. 1, pp. 1–12, 2004.
- [25] A. Ghosh, S. Achutha, W. Heidrich, and M. O'Toole, "BRDF acquisition with basis illumination," in *ICCV*, 2007.
- [26] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," ACM Trans. Graphics, vol. 22, no. 3, pp. 759–769, Jul. 2003.



Michael De Zeeuw received his B.S. in Electrical and Computer Engineering from Calvin University in 2018. He is currently pursuing his Ph.D. in ECE at Carnegie Mellon University. He received the Pittsburgh Chapter Award from the Achievement Rewards for College Scientists Foundation in 2018. His research interests include capturing and analyzing light fields, efficient shape estimation, and reflectance capture and modeling.



Aswin C. Sankaranarayanan is a professor in the ECE department at CMU, where he is the PI of the Image Science Lab. His research interests are broadly in compressive sensing, computational photography, signal processing and machine vision. His doctoral research was in the University of Maryland where his dissertation won the distinguished dissertation award from the ECE department in 2009. Aswin is the recipient of best paper awards at CVPR 2019 and ICCP 2021, the CIT Dean's Early Career

Fellowship, the Spira Teaching award, the NSF CAREER award, the Eta Kappa Nu (CMU Chapter) Excellence in Teaching award, and the Herschel Rich Invention award from Rice University.