

Axial Light Field for Curved Mirrors: Reflect Your Perspective, Widen Your View

Yuichi Taguchi, Amit Agrawal, Srikumar Ramalingam, Ashok Veeraraghavan

TR2010-038 July 2010

Abstract

Mirrors have been used to enable wide field-of-view (FOV) catadioptric imaging. The mapping between the incoming and reflected light rays depends non-linearly on the mirror shape and has been well-studied using caustics. We analyze this mapping using two-plane light field parameterization, which provides valuable insight into the geometric structure of reflected rays. Using this analysis, we study the problem of generating a single-viewpoint virtual perspective image for catadioptric systems, which is unachievable for several common configurations. Instead of minimizing distortions appearing in a single image, we propose to capture all the rays required to generate a virtual perspective by capturing a light field. We consider rotationally symmetric mirrors and show that a traditional planar light field results in significant aliasing artifacts. We propose axial light field, captured by moving the camera along the mirror rotation axis, for efficient sampling and to remove aliasing artifacts. This allows us to computationally generate wide FOV virtual perspectives using a wider class of mirrors than before, without using scene priors to depth estimation. We analyze the relationship between the axial light field parameters and the FOV/resolution of the resulting virtual perspective. Real results using a spherical mirror demonstrate generating 140 degrees FOV virtual perspective using multiple 30 degrees FOV images.

CVPR 2010

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Axial Light Field for Curved Mirrors: Reflect Your Perspective, Widen Your View

Yuichi Taguchi, Amit Agrawal, Srikumar Ramalingam, and Ashok Veeraraghavan

Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA

{taguchi, agrawal, ramalingam, veerarag} at merl.com

Abstract

Mirrors have been used to enable wide field-of-view (FOV) catadioptric imaging. The mapping between the incoming and reflected light rays depends non-linearly on the mirror shape and has been well-studied using caustics. We analyze this mapping using two-plane light field parameterization, which provides valuable insight into the geometric structure of reflected rays. Using this analysis, we study the problem of generating a single-viewpoint virtual perspective image for catadioptric systems, which is unachievable for several common configurations.

Instead of minimizing distortions appearing in a single image, we propose to capture all the rays required to generate a virtual perspective by capturing a light field. We consider rotationally symmetric mirrors and show that a traditional planar light field results in significant aliasing artifacts. We propose axial light field, captured by moving the camera along the mirror rotation axis, for efficient sampling and to remove aliasing artifacts. This allows us to computationally generate wide FOV virtual perspectives using a wider class of mirrors than before, without using scene priors or depth estimation. We analyze the relationship between the axial light field parameters and the FOV/resolution of the resulting virtual perspective. Real results using a spherical mirror demonstrate generating 140° FOV virtual perspective using multiple 30° FOV images.

1. Introduction

Image distortions from a curved mirror have been exploited to enable wide FOV catadioptric imaging. Baker and Nayar [1] presented the complete class of single-viewpoint catadioptric systems (SVCS). Theoretically, SVCS is achievable using only specific camera-mirror pairs. Some of these configurations such as a perspective camera with a sphere is not practical, since it requires the viewpoint and the center of the sphere to coincide. Previous efforts have mainly focused on removing image distortions in the resulting non-single viewpoint image, by designing mirror shapes or using scene priors/3D estimation.




Figure 1. (Left) 3 out of 25 images captured by moving a camera towards a mirror ball. (Right) Cube map output depicting a 140° FOV perspective image from a viewpoint inside the ball.

In this paper, we propose to *computationally* achieve a single-viewpoint catadioptric system, by capturing *all* the required rays to generate a virtual perspective from a viewpoint inside the mirror. We refer to these set of rays as a *virtual light field*. The virtual light field can be acquired by capturing multiple images by moving the camera. A brute force approach would be to capture the entire 4D light field (LF) by a dense sampling and choose the necessary rays to generate the virtual perspective. However, this is impractical and inefficient. Thus, it is important to understand the geometric structure of the reflected light rays from a mirror to devise an efficient sampling scheme, which reduces the number of required images.

The structure of the reflected light rays from a mirror has been well-studied using caustics/catacaustics, which describe the envelope of reflected rays or a surface tangent to all the rays. Instead, we formulate and analyze the mapping of incoming and reflected light rays for a mirror using two-plane light field parameterization, which has been extensively used to study Lambertian surfaces and their associated light fields. We show that this analysis provides novel insight into the geometry of reflected rays from a mirror. Using this analysis, we show that for rotationally symmet-

ric mirrors, the rays in the virtual light field are concentrated on certain 3D slices of the 4D light field when the virtual viewpoint is on the mirror rotation axis. Conventionally, a 4D light field is captured by placing the viewpoints on a 2D plane. We refer to it as planar light field. We show that for rotationally symmetric mirrors, a planar light field does not efficiently capture the virtual light field. We then propose *axial* light field, captured by placing the optical center of a perspective camera on the rotation axis of the mirror and moving the camera along it. The axial light field in-fact captures the same 3D slices of the 4D light field on which the information in the virtual light field is concentrated, resulting in an efficient sampling.

1.1. Contributions

- We analyze mirror reflection using two-plane light field parameterization, which provides valuable insight into the geometric structure of reflected rays.
- We show that a single-viewpoint catadioptric system can be computationally achieved by capturing all the required rays to generate a virtual perspective.
- We propose axial light field to efficiently sample the rays in the virtual light field for rotationally symmetric mirrors.
- We show that resolution properties of such a catadioptric system can be modified by varying the focal length of captured images.

1.2. Benefits and limitations

By capturing a light field instead of a single image with a mirror, wide FOV single-viewpoint perspective images can be generated for a wider class of mirrors. This enables common configurations such as a sphere/parabola with perspective cameras. Capturing a light field allows virtual perspective generation without any scene priors, 3D reconstruction, distortion correction, or non-linear optimization. In fact, the virtual viewpoint can also be varied inside the mirror to generate multiple virtual perspective images. However, axial light field requires precise camera motion along the rotation axis and the scene to be static during capture time. A multiple camera configuration similar to camera arrays is difficult since the cameras would occlude each other.

1.3. Related work

Single-viewpoint catadioptric systems: Baker and Nayar [1] presented the complete class of single-viewpoint catadioptric configurations with detailed solutions and degenerate cases. Nalwa [14] produced a $360^\circ \times 50^\circ$ FOV single-viewpoint camera using four planar mirrors in a pyramidal configuration along with four perspective cameras. SVCS can be achieved by placing the optical center of a perspective camera at the (a) center of a spherical mirror or (b) foci of an elliptical/hyperbolic mirror [17, 25], and by placing an orthographic camera parallel to the axis of a parabolic mirror. A spherical mirror with a perspective




Figure 2. (Left) Parameterization of mirror reflection in 2D. (Right) A point light source on x plane leads to a vertical line in x - u light field space. If the light source is moved by a distance z , the light field gets sheared according to z as shown in bottom.

camera is unpractical and leads to a non-single viewpoint or multi-viewpoint image (MVI) when the camera is placed outside the sphere. We show how to generate a perspective image using multiple MVI's.

Distortions in catadioptric systems: Several researches have studied image distortions in wide-angle lenses and catadioptric systems [23, 21, 20, 18]. Swaminathan *et al.* [23] proposed to use the cusp of the caustic of rotationally symmetric catadioptric systems as the virtual viewpoint to minimize distortions. In [22], authors proposed to design mirror shapes to minimize image distortions. Yu and McMillan [26] proposed a framework based on general linear cameras to assist mirror design and characterize distortions in catadioptric systems. In [21], authors derived a metric to quantify caustic distortions and presented an approach to compute minimally distorted views using priors on scene structure. It was argued that a perspective image without distortions cannot be achieved from a single MVI without knowing the scene structure. In contrast, we show that a perspective image can be generated from *multiple* MVI's without the knowledge of scene structure for rotationally symmetric catadioptric systems, by moving the camera along the rotation axis.

Axial cameras: In an axial camera, all the projection rays pass through a single line in space [15]. Cameras falling into this class include (a) stereo systems consisting of multiple perspective cameras with collinear optical centers, (b) non-central catadioptric cameras with a rotationally symmetric mirror and a pinhole camera placed on the axis of symmetry, and (c) x -slit cameras [6] and their special case of linear pushbroom cameras [9]. Our *axial light field* is similar to these axial cameras: we capture multiple perspective images by moving the camera along the rotation axis of a rotationally symmetric mirror. In contrast to a radial imaging system [11], the camera does not view the scene directly and the mirror is not hollow (scene is on same side of the mirror as the camera).

Light fields: The concept of light field as a representation of all rays of light in free-space was proposed by Levoy and Hanrahan [13] and Gortler *et al.* [8]. Chai *et al.* [3] presented frequency domain analysis for sampling and interpolation of light fields. Isaksen *et al.* [10] pro-

posed dynamic reparameterization of light fields for interactive rendering. Light fields captured using camera arrays have been used for synthetic aperture photography and digital refocusing [12, 24]. Frequency domain analysis of light transport has been presented in [5], which was extended to gradients [16] for efficient rendering. Ray-space analysis for mirrors has been studied using path perturbation for fast rendering [4]. We use similar light field analysis and show novel insight into wide FOV virtual perspective generation.

2. Light fields for curved mirrors

We first formulate mirror reflection of light field using two-plane parameterization and understand its properties in flatland. Consider a mirror surface in 2D given by $z = f(x)$ (Figure 2). We parameterize the light field by x - u with the x plane at a distance l from the mirror. Similar to [3], the u plane is set parallel to the x plane at unit distance and its coordinate is expressed in the local frame of x . Consider a ray in direction $\vec{r}_{in} = (u_{in}, -1)$ originating from x_{in} . This ray intersects the mirror surface at $f(x_0)$ with the normal vector $\vec{n} = (-f_x(x_0), 1)$. The reflected ray \vec{r}_{out} intersects the outgoing LF planes at (x_{out}, u_{out}) . Using normalized vectors ($\hat{v} = \vec{v}/|\vec{v}|$), the reflection is described as

$$\hat{r}_{out} = 2(\hat{n} \cdot \hat{r}_{in})\hat{n} - \hat{r}_{in}, \quad (1)$$

which gives

$$u_{out} = \frac{2f_x(x_0) - (1 - f_x(x_0)^2)u_{in}}{1 - f_x(x_0)^2 + 2f_x(x_0)u_{in}} \quad (2)$$

$$x_{out} = x_{in} + (l - f(x_0))(u_{in} - u_{out}). \quad (3)$$

Note that the above mapping includes the unknown point of reflection x_0 , which needs to be solved for using the following constraint: $x_0 = x_{in} + (l - f(x_0))u_{in}$.

2.1. LF visualization for common configurations

We first consider reflected light fields due to real scene points outside the mirror. Consider a point source at x_{in} , emitting rays in different directions. The incoming light field is a vertical line in x - u space. For a point light source on a plane at distance z from the x plane, the light field gets sheared [5] according to z (Figure 2). This incoming light field gets reflected by the mirror.

For a flat mirror, $f(x) = f_x(x) = 0$ for all x . Thus, $u_{out} = -u_{in}$ and $x_{out} = x_{in} + 2lu_{in}$. The incoming light field therefore gets sheared (Figure 3(a)), showing that flat mirror shifts the viewpoint. In contrast, for a curved mirror, the reflection of a scene point traces a curve in the output light field, as shown in Figure 3(a) for a sphere.

Light rays for a virtual viewpoint: Now consider all the rays reflected from a radiant point inside the mirror. This scenario is similar to having a virtual viewpoint inside the mirror. All the rays which would pass through the virtual viewpoint after getting reflected by the mirror should be captured to generate a virtual perspective. We refer to these set of rays as a virtual light field. A SVCS enforces the




Figure 3. (a) Reflected light fields, (b) planar LF sampling and (c) axial LF sampling in flatland.




Figure 4. The set of all rays which pass through the virtual viewpoint (within a FOV) after reflecting by the mirror is termed as virtual LF $L(x, y, u, v)$. Each camera in a planar LF captures a u - v slice P at spatial location (n, m) . In an axial LF, each camera captures a sheared 2D slice Q depending on its distance z .

constraint that all these reflected rays should pass through a *single* viewpoint outside the mirror, at which a pinhole camera is placed. Thus, it enforces the virtual light field to be a *line* in x - u space. As expected, this is a hard geometric constraint. Figure 3(a) shows the virtual light field for a virtual viewpoint inside a sphere. In general, the virtual light field for a curved mirror will form a curve, which cannot be captured by a single viewpoint. This curve, however, could be sampled by capturing multiple images.

2.2. Capturing the virtual light field

Let $L(x, u)$ be the virtual light field in flatland. In a planar light field, the viewpoint is changed along the x axis. Thus, a camera placed at $x = x_i$ produces an image $i(u) = L(x_i, u)$, which samples a *vertical* slice of the light field (Figure 3(b)). In an axial light field, the viewpoint is changed perpendicular to the x axis. Here, each image corresponds to a sheared slice $i(u) = L(uz_i, u)$, where z_i is the distance of the viewpoint from the x plane (Figure 3(c)). Next, we analyze the sampling in 4D and show that the axial sampling is aligned with the structure in the virtual light field for rotationally symmetric mirrors.

3. Light field based SVCS

As explained earlier, a traditional SVCS requires all rays in the virtual light field to pass through a single viewpoint. This restricts the mirror-camera configuration to those presented in [1]. Theoretically, by capturing multiple images (light field), single-viewpoint virtual perspectives can be obtained for any mirror by sampling all the required rays. This in practice would be inefficient and for a general mirror, such a sampling pattern will not have any regular structure. Our goal is to achieve a middle ground, i.e., significantly enlarge the family of mirrors, while keeping the sampling overhead minimal.




Figure 5. Planar LF sampling. (a) x-u-v slices of virtual LF and 2D camera slices (planes) at various x for $y = 0$. Only the camera on the axis captures a cone of rays. (b) Similar plot for $y > 0$ shows that information in other slices is sparse. (c) A planar LF samples vertical segments in x-u space.




Figure 6. Axial LF sampling. (a) x-u-v slices of virtual LF and 2D camera slices (planes) for various z . The camera planes align with the circle in each slice and thus each camera captures a cone of rays. (b) Information in virtual LF is along slanted lines in x-u space, which is captured by moving the camera along the rotation axis. (c) Each camera in axial LF captures a cone of rays, which contributes to a circle of pixels in the virtual perspective image.

Consider the class of mirrors which are *rotationally symmetric*. These mirrors can be obtained by rotating a 2D curve along an axis. For such mirrors, if the virtual viewpoint is on the rotation axis, all the rays in the virtual light field pass through a *line* (the rotation axis). Thus, this configuration relaxes the fixed viewpoint constraint in SVCS to a line. The virtual light field now has structure, which can be exploited for efficient sampling. We show that for rotationally symmetric mirrors, an axial light field provides efficient sampling of the virtual light field. Such a sampling allows generation of perspective images from virtual viewpoints on the rotation axis inside the mirror, *without* any scene priors, 3D estimation, or distortion correction.

3.1. Rotationally symmetric mirrors

To understand such sampling, consider a sphere of radius r centered at origin and a desired virtual viewpoint at distance d inside the mirror. Consider LF planes at a finite distance as shown in Figure 4. Let $L(x, y, u, v)$ be the 4D virtual LF. Note that the virtual LF only contains the rays that pass through the virtual viewpoint, not all the rays. While catacaustic with radiant point inside the mirror describes the *envelope* of all the reflected rays, the virtual LF describes the *intersection* of reflected rays with LF planes.

Planar LF sampling: A camera located at $x = n$, $y = m$ on the xy LF plane captures a planar 2D slice $P(u, v) = L(n, m, u, v)$ of the virtual LF. Since it is not possible to show the resulting LF in 4D, Figure 5(a) shows the x-u-v slice of $L(x, y, u, v)$ at $y = 0$. Figure 5(a) also shows the rays captured by each camera (shown as planes)

at various x locations for $y = 0$. Note that only when the location of the camera coincides with the axis, the camera captures a cone of rays as a *circle* in the image plane. Otherwise, it captures few rays (a maximum of three rays for sphere as derived in supplementary materials). Figure 5(b) shows a different 3D slice $L(x, y = 5, u, v)$ for non-zero y . Notice that this slice contains only a few rays (denoted by \times). Thus, the information in slices $L(x, y, :, :)$ of the virtual LF is sparse, unless x, y lie on the mirror axis. Next we show that information is concentrated in different kind of slices and an axial LF captures it.

Axial LF sampling: Consider a camera located on the axis at a distance z from the xy LF plane. This camera captures a planar 2D slice $Q(u, v) = L(uz, vz, u, v)$ of the virtual LF. Notice that this 2D slice (image) is different from the one in the planar LF case. Figure 6(a) plots the slices $L(uz, vz, u, v)$ of the virtual LF for several z along with the rays captured by each camera (shown as planes). Notice that for each slice of $L(uz, vz, u, v)$, the central cone of rays can be captured by the corresponding camera. The x-u slices of $L(x, y, u, v)$ show that the information in virtual LF is concentrated along lines passing through the center (axis). Each camera in a planar LF captures a vertical segment in x-u slice. In contrast, each camera in an axial LF captures a tilted line segment, which is in-fact aligned with the structures in the virtual LF and is a better sampling strategy. The u-v slices show how each camera captures a cone of rays as a circle in the image plane.

In summary, the information in the virtual LF for a rotationally symmetric mirror is concentrated in a 3D subset of rays $L(uz, vz, u, v)$, which is captured efficiently by an axial LF. In contrast, a planar LF captures a 4D subset of rays. Thus, an axial LF provides better sampling than a planar LF to generate a virtual perspective using the same number of images.

Interpretation using caustics: A caustic [2, 20] represents a surface that is tangent to all the rays passing from a real viewpoint and reflected by the mirror. Only if the caustic is a point, the camera+mirror system becomes a SVCS. Thus, the *common* rays shared by a real camera and a virtual camera are those that pass through the virtual viewpoint and are tangent to the caustic. Since the incoming ray, normal of the reflection point, and the reflected ray should lie on a single plane, the rotation axis and real and virtual viewpoints define a plane of reflection (Figure 7). Note that only the rays on this plane that are tangent to the caustic can be captured at the real viewpoint. When the virtual and real viewpoints are both on the rotation axis, the plane of reflection degenerates into the mirror rotation axis. Thus, there is an infinite family of planes that contain the rotation axis, and all rays on all these planes (a cone of rays) can be captured at the real viewpoint. Since in axial LF, all the real viewpoints lie on the rotation axis and capture a cone of rays, it is more efficient than planar LF sampling.




Figure 7. When the real and virtual viewpoints are on the rotation axis, all rays on an infinite family of planes (a cone of rays) can be captured. Otherwise, only those rays that lie on a single reflection plane are captured by the real camera.

4. Virtual perspective image from axial LF

We derive the parameters of the axial LF for a sphere of radius r and virtual viewpoint at a distance d inside the sphere (Figure 8(a)). Let θ_v be the angle of ray from the virtual viewpoint, intersecting the sphere at an angle ϕ as shown. Then,

$$\tan \theta_v = \frac{r \cos \phi}{r \sin \phi - r + d}. \quad (4)$$

After reflection, the ray passes through $(0, r + z)$, which gives

$$r + z - r \sin \phi = -r \cos \phi \tan(2\phi - \alpha) \quad (5)$$

$$\sin \phi = \frac{r(z + 2r - d)}{2(r - d)(z + r)}. \quad (6)$$

For a given θ_v , one can compute ϕ using (4) and z using (6). If a real perspective camera is placed at distance z from the top of the sphere, then a cone of rays corresponding to θ_v can be captured as a cone of rays corresponding to θ_r , where $\theta_r = \arctan(\frac{r \cos \phi}{z + r - r \sin \phi})$. Figure 9 depicts how to generate a virtual perspective image from axial LF images.

4.1. Axial LF viewpoints

For a given virtual FOV, the positions at which the camera should be placed along the axis can be obtained from the above equations (for sphere). Figure 10(a) plots the axial LF viewpoints with respect to θ_v for several rotationally symmetric mirrors. Each mirror was generated by revolving a 2D curve along the rotation axis, which is shown in Figure 8(c). For mirrors other than sphere, θ_r and z were computed numerically for each θ_v . Notice that for a sphere, the virtual viewpoint location increases rapidly after 140° FOV compared to a parabola. For a conical mirror, smaller θ_v cannot be captured since it leads to negative z , which is physically impossible.

4.2. Angular samples of axial LF

Figure 10(b) plots θ_r with respect to θ_v for the above mirrors. Note that each θ_r is associated with corresponding z shown in Figure 10(a). Some interesting observations can be made here. For a sphere, the maximum FOV required in any input image is less than 20° and an output FOV of

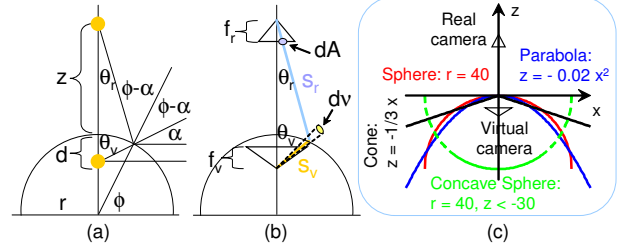


Figure 8. (a) Determining z and θ_r for a given θ_v using a sphere of radius r and virtual viewpoint at d . (b) The resolution of virtual camera is s_v^2/s_r^2 times that of real camera [1]. (c) Mirror parameters used in Figure 10 plots and Figure 11 simulations.




Figure 9. (Left) For a rotationally symmetric mirror, the virtual perspective is obtained by copying a resized circle of pixels from each image. (Right) Ray interpolation using a virtual depth plane produces better results than using infinite depth assumption.

140° can be easily achieved. However, note that θ_r starts decreasing after a certain extent as θ_v increases. This implies that large angles (circles) in virtual image are obtained from smaller angles (circles) as θ_v increases. Therefore, it would lead to blurring of the virtual image in periphery and z dependent angular resolution. We analyze the resolution properties in detail in Section 4.3. A parabolic mirror reduces blurring compared to sphere since θ_r does not decrease with increasing θ_v . Interestingly, for a concave sphere, θ_r increases with θ_v and thus it would provide even better angular resolution in periphery. However, concave mirrors restrict the virtual FOV, because θ_r keeps increasing and exceeds the FOV of the real camera. The simulations shown in Section 5 confirm these observations.

4.3. Resolution analysis

Consider an infinitesimal area dA on the image plane, which images an infinitesimal solid angle $d\nu$ of the scene. The resolution of an imaging system is then defined as $\frac{dA}{d\nu}$. For a perspective camera with focal length f_r , the resolution equals $\frac{f_r^2}{\cos^3 \theta_r}$, where θ_r is the angle between the optical axis and line joining the optical center to the center of dA . For a catadioptric system, the resolution at angle θ_v is a factor $\frac{s_v^2}{s_r^2}$ of the real camera resolution [1], where s_v and s_r are the length of rays as shown in Figure 8(b). For a sphere, we analytically derive the factor as follows:

$$\frac{s_v^2}{s_r^2} = \frac{(r \cos \phi)^2 + (r \sin \phi - r + d)^2}{(r \cos \phi)^2 + (z + r - r \sin \phi)^2} = \frac{(r - d)^2}{(z + r)^2}. \quad (7)$$

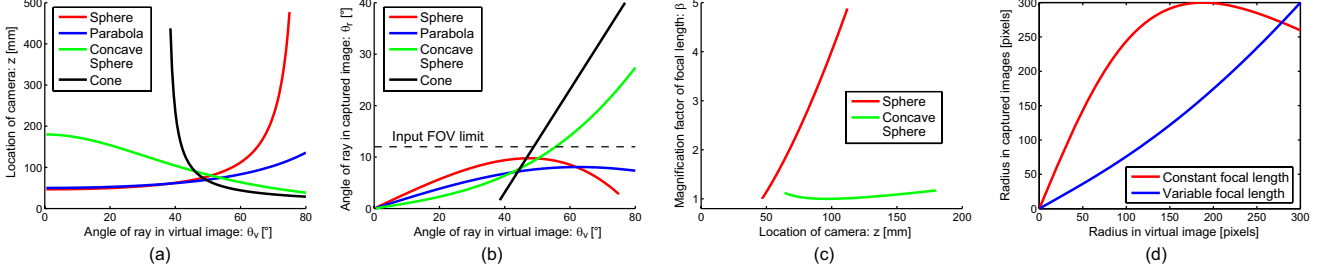


Figure 10. (a) Plot of axial LF viewpoint z for each θ_v for various mirrors. (b) Corresponding plot of θ_r . Note that θ_v and θ_r are half-angle of FOV. The horizontal line depicts the FOV of input images and restricts the FOV of virtual image. (c) Change in focal length β required to maintain angular resolution similar to a perspective camera. (d) By changing the focal length of the captured images, input circle sizes can be modified to change virtual image resolution characteristics.

Note that as θ_v changes, the position of the real camera z at which the corresponding cone of rays is captured changes, thereby changing the resolution.

Typically, the resolution of a catadioptric system can be changed only by changing the mirror shape. Mirror shapes have been proposed to achieve constant resolution in space or angle [7]. However, capturing multiple images gives an extra degree of flexibility, since each image can have different camera parameters. The angular resolution of virtual image can be made similar to a conventional perspective image (not a constant angular resolution) by varying its resolution according to $(kf_v)^2/\cos^3\theta_v$, where k is the change in focal length to achieve wider FOV for the same number of pixels. Thus, for all θ_v ,

$$\frac{(kf_v)^2}{\cos^3\theta_v} = \frac{s_v^2}{s_r^2} \frac{f_r^2}{\cos^3\theta_r} = \frac{(r-d)^2}{(z+r)^2} \frac{f_r^2}{\cos^3\theta_r} \quad (8)$$

should be satisfied. We define

$$\beta(z) = \frac{f_r}{kf_v} = \frac{(z+r)\cos^{3/2}\theta_r}{(r-d)\cos^{3/2}\theta_v} \quad (9)$$

as the required change in the focal length of the real camera with respect to the viewpoint z . Note that $\beta(z) = 1$ implies same focal length f_r for all images. Thus, the axial LF should be captured by moving the camera away from the sphere while increasing its focal length.

Effect in image space: Note that the virtual image is obtained by copying a *resized* circle of pixels from each captured image. The circle size and its resizing depends on θ_r and θ_v . For viewpoints away from the sphere, the *inner* circle of the captured image is resized to the *outer* circle of the virtual image. This explains the blurring for larger θ_v in the virtual image. This is also evident from Figure 10(b) as θ_r decreased with θ_v for large θ_v , when focal length f_r is not changed. By changing f_r based on (9), the size of each captured circle can be made larger to achieve a desired resolution as shown in Figure 10(d).

Concave mirrors reverse the above relationship (θ_r does not decrease as θ_v increases). Thus, concave mirrors offer less blurring in image periphery. Figure 10(c) shows that the required change in focal length β is small for a concave

sphere compared to a convex sphere. For convex sphere, β could become physically impractical for large θ_v .

5. Simulations and real results

Simulations: We simulated axial LF images for several mirrors using POV-Ray software. The mirror is placed inside a cube whose walls are painted with checkerboard of different colors. The axial LF was generated by moving the camera non-uniformly in z to maintain equal increments in θ_v . Figure 11 shows a near and a far image for each axial LF, along with the generated virtual perspective image. In each case, the virtual image was generated in the same manner as depicted in Figure 9, with z and θ_r depending on mirror shape for each θ_v . We did not use any image/scene priors, non-linear optimization, or distortion correction. While the captured images show severe distortions, the resulting virtual image is close to ground truth (shown in Figure 12). All straight lines on all scene planes are mapped as straight lines in the image. As discussed previously, for a sphere, the virtual image has more blur in periphery compared to a parabola and concave mirror.

Ray interpolation: In practice, the camera cannot be placed at all z corresponding to each θ_v . This requires to interpolate rays using images from neighboring axial viewpoints. In traditional LF, appropriate ray interpolation is performed by assuming scene depths to be within z_{min} and z_{max} and defining a constant depth plane at z_{opt} , where $\frac{1}{z_{opt}} = (\frac{1}{z_{min}} + \frac{1}{z_{max}})/2$ [3]. For a spherical mirror, all the scene depths lie within a *virtual depth range* inside the mirror [19]. We therefore use a virtual depth plane passing through the virtual viewpoint, which is a good approximation of the constant depth plane, and compute rays using bilinear interpolation (Figure 9). This gives better results than using an infinite depth assumption as shown in Figure 12.

5.1. Real results

We demonstrate our technique by generating 140° FOV virtual perspective images from inside a sphere using an axial LF. We use a stainless steel mirror ball of radius 38.1 mm, imaged with a Pointgrey Dragonfly camera mounted on a robot arm that allows precise motion. The camera uses

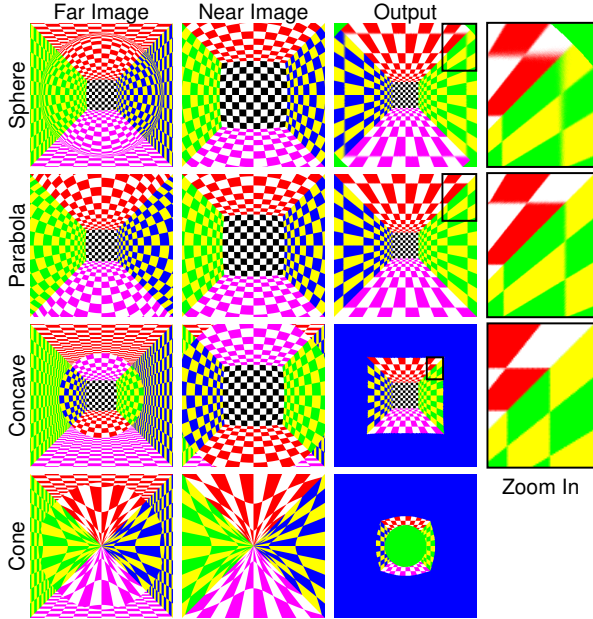


Figure 11. Simulation results using POV-ray for various mirror shapes assuming the same virtual FOV of 140° . The green and blue regions depict rays which cannot be captured due to limited z positions and limited FOV of the input camera, respectively.

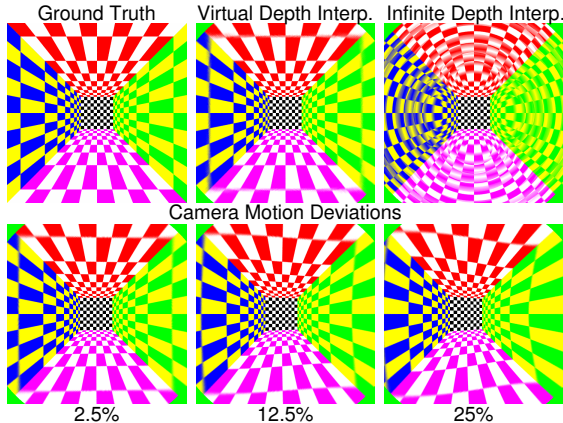


Figure 12. (Top) Ground truth for Figure 11 and virtual images for sphere using different interpolation methods. (Bottom) Incorrect estimation of the center of sphere as a percentage of sphere radius along x axis (perpendicular to camera motion direction) leads to increased distortions.

an 8 mm computer lens, giving an input FOV of $32^\circ \times 24^\circ$.

Calibration: The intrinsic camera parameters are computed off-line using checkerboard images. The mirror ball is placed on a checkerboard. We compute the pose of the camera with respect to the checkerboard and move the camera to make the sensor plane parallel to the checkerboard. We then manually mark the sphere boundary in the image to determine rays tangential to the sphere and their central ray, which makes the same angle with all the tangential rays. Using the fact that the sphere center lies on the central ray at a distance of r from the checkerboard, its location can

be found in the camera coordinate system. The camera is then positioned so that the sphere center lies on the optical axis. The sphere center is estimated again to improve accuracy. Using the camera and the sphere centers, the direction vector is obtained along which the camera is moved.

Figure 1 shows 3 out of 25 captured images. The distance z of the camera was changed non-uniformly from 41 mm to 334 mm using the θ_v - z relationship for equal increments in θ_v . The scene consists of several checkerboards on ceiling, two books, a color Macbeth chart, a toy and a textured board on the sides close to the mirror ball. Figure 1 shows the synthesized cube map from a virtual viewpoint 13 mm inside the sphere, made by joining 5 images (center image has $90^\circ \times 90^\circ$ FOV, while others have $90^\circ \times 45^\circ$ FOV). Notice that straight lines in both orthogonal directions in all (both near and far) scene planes are maintained as straight lines in the virtual image. Figure 13 compares real results on a different scene using axial and planar LF sampling. For planar LF, we captured 5×5 images at a constant height of 200 mm with a baseline interval of 13 mm. Blurring and aliasing artifacts can be seen clearly in the planar LF result because of the inefficient sampling, while our axial LF sampling generates a sharp result.

6. Discussions and conclusions

Off-the-shelf CCD cameras typically have more FOV in x direction than y direction due to non-square sensor size. However, the virtual perspective image will have the same FOV in x and y directions due to mirror being rotationally symmetric. Thus, at large enough angles, some rays in the vertical direction may not be sampled although other rays within the same cone in the horizontal direction can be sampled. We assume that all captured images are in sharp focus, but out-of-focus regions will blur the final image at corresponding regions. We currently use a small aperture to get sufficient depth of field in captured images. If the camera motion is not along the rotation axis, it results in distortions and blurring as shown in Figure 12. Non-linear processing and higher level information such as ego-motion/depth estimation may be utilized to remove such distortions and blurring. While a complete frequency domain analysis of light fields for general mirrors (similar to [3]) is beyond the scope of the paper, it would provide further insights into optimal sampling for virtual perspective generation.

Conclusions: Common mirror configurations with perspective cameras do not lead to a single-viewpoint catadioptric system. Instead of minimizing distortions in the resulting image, we propose to capture all the required rays to generate a virtual perspective image by moving the camera. We analyzed these set of rays in terms of a virtual light field. We showed that traditional planar light field does not result in efficient sampling, but the proposed axial light field samples the rays in the virtual light field efficiently for rotationally symmetric mirrors. Real results using a mirror ball




Figure 13. Comparison between planar and axial LF sampling using the same number of input images (25) and the same FOV of virtual perspective. Notice the blur and aliasing artifacts in planar LF sampling on flower and mannequin. In both cases, the camera is not correctly synthesized since it moves during capture.

demonstrates that wide FOV virtual perspective can be generated from axial light field without using any scene priors or distortion correction. We presented an extensive analysis relating axial light field parameters and FOV/resolution characteristics of the resulting virtual image.

Acknowledgments: We thank the anonymous reviewers for their suggestions. We also thank Jay Thornton, Keisuke Kojima, John Barnwell, and Haruhisa Okuda, Mitsubishi Electric, Japan, for help and support.

References

- [1] S. Baker and S. Nayar. A theory of single-viewpoint catadioptric image formation. *Int'l J. Computer Vision*, 35(2):175–196, Nov 1999.
- [2] M. Born and E. Wolf. *Principles of Optics*. Pergamon Press, 1965.
- [3] J.-X. Chai, S.-C. Chan, H.-Y. Shum, and X. Tong. Plenoptic sampling. In *SIGGRAPH*, pages 307–318, 2000.
- [4] M. Chen and J. Arvo. Theory and application of specular path perturbation. *ACM Trans. Graph.*, 19(4):246–278, 2000.
- [5] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. X. Sillion. A frequency analysis of light transport. *ACM Trans. Graph.*, 24(3):1115–1126, 2005.
- [6] D. Feldman, T. Pajdla, and D. Weinshall. On the epipolar geometry of the crossed-slits projection. In *ICCV*, 2003.
- [7] J. Gaspar, C. Decco, J. Okamoto, and J. Santos-Victor. Constant resolution omnidirectional cameras. In *OMNIVIS*, pages 27–34, 2002.
- [8] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The lumigraph. In *SIGGRAPH*, pages 43–54, 1996.
- [9] R. Gupta and R. Hartley. Linear pushbroom cameras. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(9):963–975, 1997.
- [10] A. Isaksen, L. McMillan, and S. Gortler. Dynamically reparameterized light fields. In *SIGGRAPH*, pages 297–306, 2000.
- [11] S. Kuthirummal and S. Nayar. Multiview radial catadioptric imaging for scene capture. *ACM Trans. Graph.*, 25(3):916–923, 2006.
- [12] M. Levoy, B. Chen, V. Vaish, M. Horowitz, M. McDowall, and M. Bolas. Synthetic aperture confocal imaging. *ACM Trans. Graph.*, 23:825–834, 2004.
- [13] M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH*, pages 31–42, 1996.
- [14] V. Nalwa. A true omnidirectional viewer. In *Technical report, Bell Laboratories, Holmdel, NJ, USA*, 1996.
- [15] S. Ramalingam, P. Sturm, and S. Lodha. Theory and calibration algorithms for axial cameras. In *ACCV*, 2006.
- [16] R. Ramamoorthi, D. Mahajan, and P. Belhumeur. A first-order analysis of lighting, shading, and shadows. *ACM Trans. Graph.*, 26(1), Jan. 2007.
- [17] D. W. Rees. Panoramic television viewing system. In *United States Patent (3,505,465)*, 1970.
- [18] S. Shah and J. Aggarwal. Intrinsic parameter calibration procedure for a (high distortion) fish-eye lens camera with distortion model and accuracy estimation. In *Pattern Recognition*, volume 29, pages 1775–1788, 1996.
- [19] R. Swaminathan. Focus in catadioptric imaging systems. In *ICCV*, pages 1–7, 2007.
- [20] R. Swaminathan, M. Grossberg, and S. Nayar. Caustics of catadioptric cameras. In *ICCV*, volume 2, pages 2–9, 2001.
- [21] R. Swaminathan, M. Grossberg, and S. Nayar. A perspective on distortions. In *CVPR*, volume 2, pages 594–601, 2003.
- [22] R. Swaminathan, M. Grossberg, and S. Nayar. Designing mirrors for catadioptric systems that minimize image errors. In *OMNIVIS*, 2004.
- [23] R. Swaminathan, M. Grossberg, and S. Nayar. Non-single viewpoint catadioptric cameras: Geometry and analysis. *Int'l J. Computer Vision*, 66(3):211–229, Mar 2006.
- [24] V. Vaish, B. Wilburn, N. Joshi, and M. Levoy. Using plane + parallax for calibrating dense camera arrays. In *CVPR*, pages 2–9, 2004.
- [25] K. Yamazawa, Y. Yagi, and M. Yachida. Obstacle avoidance with omnidirectional image sensor hyperomni vision. In *ICRA*, pages 1062–1067, 1995.
- [26] J. Yu and L. McMillan. Modelling reflections via multiperspective imaging. In *CVPR*, volume 1, pages 117–124, 2005.