

A General Imaging Model and a Method for Finding its Parameters

Michael D. Grossberg and Shree K. Nayar

Department of Computer Science, Columbia University
New York, New York 10027
{mdog, nayar}@cs.columbia.edu

Abstract

Linear perspective projection has served as the dominant imaging model in computer vision. Recent developments in image sensing make the perspective model highly restrictive. This paper presents a general imaging model that can be used to represent an arbitrary imaging system. It is observed that all imaging systems perform a mapping from incoming scene rays to photo-sensitive elements on the image detector. This mapping can be conveniently described using a set of virtual sensing elements called raxels. Raxels include geometric, radiometric and optical properties. We present a novel calibration method that uses structured light patterns to extract the raxel parameters of an arbitrary imaging system. Experimental results for perspective as well as non-perspective imaging systems are included.

1 Introduction

Since the Renaissance, artists have been fascinated by the visual manifestations of perspective projection. Geometers have studied the properties of the pinhole imaging model and derived a large suite of projective invariants that provide insights into the relation between a scene and its perspective image. The field of optics has developed high quality imaging lenses that closely adhere to the perspective model. Today, perspective projection serves as the dominant imaging model in computer vision and computer graphics.

Despite its great relevance, there are several reasons that make the perspective model far too restrictive. In recent years, the notion of a “vision sensor” has taken on a much broader meaning. A variety of devices have been developed that sample the light field [8] or the plenoptic function [1] associated with a scene in interesting and useful non-perspective ways. Figure 1 shows some examples of such imaging systems that are widely used today. Figure 1(a) shows a catadioptric sensor that uses a combination of lenses and mirrors. Even when such a sensor has a single effective viewpoint [21], its projection model can include a variety of non-perspective distortions (barrel, pincushion, or more complex)[2]. More interesting is the fact that certain applications (see [22] for example) require the system to not have a single viewpoint but rather a locus of viewpoints (catacaustic [4]). Similarly, wide-angle lens systems [20]

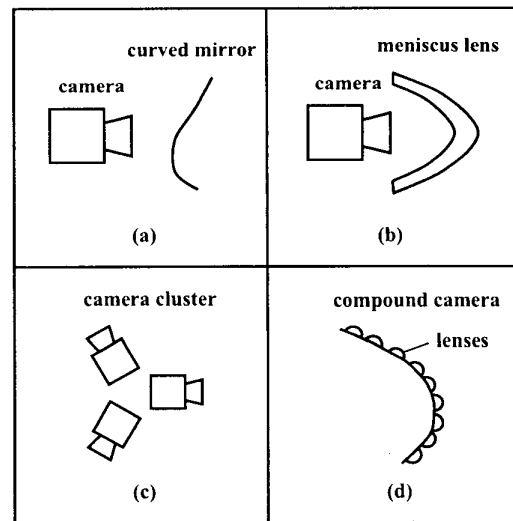


Figure 1: Examples of non-perspective imaging systems: (a) a catadioptric system, (b) a dioptic wide-angle system, (c) an imaging system made of a camera cluster, and (d) a compound camera made of individual sensing elements, each including a receptor and a lens. In all of these cases, the imaging model of the system deviates from perspective projection.

like the one shown in Figure 1(b), include severe projective distortions and often have a locus of viewpoints (called a dicaustic). Recently, clusters of cameras, like the one shown in Figure 1(c), have become popular [17][24]. It is clear that such a system includes multiple viewpoints or loci of viewpoints, each one associated with one of the cameras in the cluster. Finally, in the case of insects, nature has evolved eyes that have compound lenses [7], [6], such as the one shown in Figure 1(d). These eyes are composed of thousands of “ommatidia”, each ommatidium including a receptor and lens. It is only a matter of time before we see solid-state cameras with flexible imaging surfaces that include a large number of such ommatidia.

In this paper we address two questions that we believe are fundamental to imaging:

- Is there an imaging model that is general enough to represent any arbitrary imaging system? Note that we

are not placing any restrictions on the properties of the imaging system. It could be perspective or non-perspective.

- Given an unknown imaging system (a black box), is there a simple calibration method that can compute the parameters of the imaging model? Note that for the types of systems we wish to subsume in our imaging model, conventional camera calibration techniques will not suffice.

Such a general imaging model must be flexible enough to cover the wide range of devices that are of interest to us. Yet, it should be specific enough in terms of its parameters that it is useful in practice. Our approach is to exploit the fact that all imaging systems perform a mapping from incoming scene rays to photo-sensitive elements on the image detector. This mapping can be conveniently described by a *ray surface* which is a surface in three-dimensional space from which the rays are measured in various directions. The smallest element of our imaging model is a virtual photo-sensitive element that measures light in essentially a single direction. We refer to these virtual elements as ray pixels, or *raxels*. It turns out that a convenient way to represent the ray surface on which the raxels reside is the *caustic* of the imaging system. In addition to its geometric parameters, each raxel has its own radiometric response function and local point spread function.

After describing the general imaging model and its properties, we present a simple method for finding the parameters of the model for any arbitrary imaging system. It is important to note that, given the non-perspective nature of a general device, conventional calibration methods based on known scene points [25] or self-calibration techniques that use unknown scene points [5], [10], [15], cannot be directly applied. Since we are interested in the mapping from rays to image points, we need a *ray-based* calibration method. We describe a simple and effective ray-based approach that uses structured light patterns. This method allows a user to obtain the geometric, radiometric, and optical parameters of an arbitrarily complex imaging system in a matter of minutes.

2 General Imaging Model: Geometry

We will consider the imaging system shown in Figure 2 when formulating our mathematical model of the imaging sensors mentioned in section 1. The system includes a detector with a large number of photo-sensitive elements (pixels). The detector could be an electronic chip, film, or any other light sensitive device. The technology used to implement it is not important. The imaging optics typically include several elements. Even a relatively simple optical component has about five individual lenses within it. In our arbitrary system, there may be additional optical elements such as mirrors, prisms, or beam-splitters. In fact, the system

could be comprised of multiple individual imaging systems, each with its own imaging optics and image detector.

Irrespective of its specific design, the purpose of an imaging system is to map incoming rays of light from the scene onto pixels on the detector. Each pixel collects light energy from a bundle of closely packed rays in any system that has a non-zero aperture size. However, the bundle can be represented by a single chief (or principle) ray when studying the geometric properties of the imaging system. As shown in Figure 2, the system maps the ray P_i to the pixel i . The path that incoming ray traverses to the pixel can be arbitrarily complex.




Figure 2: An imaging system directs incoming light rays to its photo-sensitive elements (pixels). Each pixel collects light from a bundle of rays that pass through the finite aperture of the imaging system. However, we will assume there is a correspondence between each individual detector element i and a specific ray of light P_i , entering the system.

If the imaging system is perspective, all the incoming light rays are projected directly onto the detector plane through a single point, namely, the effective pinhole of the perspective system. This is not true in an arbitrary system. For instance, it is clear from Figure 2 that the captured rays do not meet at a single effective viewpoint. The goal of this section is to present a geometrical model that can represent such imaging systems.

2.1 Raxels

It is convenient to represent the mapping from scene rays to pixels in a form that easily lends itself to manipulation and analysis. We can replace our physical pixels with an abstract mathematical equivalent we refer to as a ray pixel or a “raxel”. A raxel is a virtual photo-sensitive element that measures the light energy of a compact bundle of rays which can be represented as a single principle incoming ray. A similar abstraction, the pencigraph, was proposed by Mann

[16] for the perspective case. The abstract optical model of our virtual raxel is shown in Figure 3. Each raxel includes a pixel that measures light energy and imaging optics (a lens) that collects the bundle of rays around an incoming ray. In this section, we will focus on the geometric properties (locations and orientations) of raxels. However, each raxel can possess its own radiometric (brightness and wavelength) response as well as optical (point spread) properties. These non-geometric properties will be discussed in subsequent sections.




Figure 3: (a) A raxel is a virtual replacement for a real photo-sensitive element. It may be placed along the line of a principle ray of light entering the imaging system. In addition to location and orientation, a raxel may have radiometric and optical parameters. (b) The notation for a raxel used in this paper.

2.2 Plenoptic Function

What does an imaging system see? The input to the system is the plenoptic function [1]. The plenoptic function $\Phi(\mathbf{p}, \mathbf{q}, t, \lambda)$ gives the intensity of light at each point \mathbf{p} in space, from direction \mathbf{q} , at an instant of time t and wavelength λ . We specify position by (p_x, p_y, p_z) and direction by two angles (q_ϕ, q_θ) . Still images represent an integration of light energy over a short time period, given by the effective shutter speed. Further, each photo-sensitive element will average the plenoptic function across a range of wavelengths. Thus, we set aside time and wavelength by considering monochromatic still imaging. We will only consider the plenoptic function as a function of position and direction: $\Phi(\mathbf{p}, \mathbf{q})$.

We may view a raxel as a delta function¹ $\delta_{\mathbf{p}_0, \mathbf{q}_0}$ over \mathbf{p}, \mathbf{q} space, as it measures the value of plenoptic function $\Phi(\mathbf{p}, \mathbf{q})$ at $(\mathbf{p}_0, \mathbf{q}_0)$. Hence the parameters for a raxel are just position \mathbf{p}_0 and direction \mathbf{q}_0 .

2.3 Pencils of Rays and Ray Surfaces

From *where* does the system (with its raxels) see the plenoptic function? Each point in the image corresponds to a ray. Thus, the set of positions and directions determined by the set of rays is the part of the domain of the plenoptic function relevant to our system.²

¹If the raxel has a non-linear radiometric response then the response must be linearized for the raxel to be a delta function.

²The related problem of representing the positions and directions corresponding to a light source was explored in [13].

The most general imaging system is described by a list of these rays. For clarity, we assume our image is two-dimensional.³ An image point is specified by (x, y) . A scene point \mathbf{p} imaged at (x, y) can be anywhere along the corresponding ray. To specify the point in space we define a parameter r along the ray. In the perspective case, r may be chosen as scene depth.

A point $\mathbf{p}(x, y, r)$ imaged at (x, y) at depth r is imaged along a ray in the direction $\mathbf{q}(x, y, r)$. Thus, we see the plenoptic function only from those points in the range of \mathbf{p} and \mathbf{q} . We may place a raxel anywhere along a ray.⁴ It will be more convenient for representing the model to arrange the raxels on a surface we call a *ray surface*. For example, consider a sphere enclosing our imaging system, as shown in Figure 4. For each photo-sensitive element i there is some point \mathbf{p}_i on the sphere that received a ray in the direction \mathbf{q}_i . Thus we can place our raxels on the sphere by assigning them the positions and directions $(\mathbf{p}_i, \mathbf{q}_i)$. It is important to note that there could be several rays that enter the sphere at the same point but with different directions (see $\mathbf{q}_1, \mathbf{q}_2$ and \mathbf{q}_3 in Figure 4). Thus the direction \mathbf{q} is not, in general, a function of \mathbf{p} .




Figure 4: An imaging system may be modeled as a set of raxels on a sphere surrounding the imaging system. Each raxel i has a position \mathbf{p}_i on the sphere, and an orientation \mathbf{q}_i aligned with an incoming ray. Multiple raxels may be located at the same point ($\mathbf{p}_1 = \mathbf{p}_2 = \mathbf{p}_3$), but have different directions.

The choice of intersecting the incoming rays with a refer-

³Many of the arrays of photo-sensitive elements in the imaging devices described in section 1 are one or two-dimensional. Multi-camera systems can be represented by two-dimensional arrays parameterized by an extra parameter.

⁴Intensities usually do not change much along a ray (particularly when the medium is air) provided the displacement is small with respect to the total length of the ray.

ence sphere is arbitrary. In [9] and [14], it was suggested that the plenoptic function could also be restricted to a plane. The important thing is to choose some reference surface so that each incoming ray intersects this surface at only one point. If the incoming rays are parameterized by image coordinates (x, y) , each ray will intersect a reference surface at one point $\mathbf{p}(x, y)$. We can write the ray surface as a function of (x, y) as:

$$\mathbf{s}(x, y) = (\mathbf{p}(x, y), \mathbf{q}(x, y)). \quad (1)$$

We can express the position of a point along the ray as $\mathbf{p}(x, y, r) = \mathbf{p}(x, y) + r\mathbf{q}(x, y)$. This allows us to express the relevant subset of the domain of the plenoptic function as the range of

$$\mathcal{L}(x, y, r) = (\mathbf{p} + r\mathbf{q}, \mathbf{q}). \quad (2)$$

In the case of an unknown imaging system we may measure $\mathbf{s}(x, y)$ along some ray surface. In the case of a known imaging system we are able to compute $\mathbf{s}(x, y)$ *a priori*. Using equation 2 we may express one ray surface in terms of another ray surface such as a sphere or a plane.

3 Caustics

While convenient, the choices of a plane or sphere as the reference surface come with their drawbacks. The direction $\mathbf{q}(x, y)$ of a raxel need not have any relation to the position $\mathbf{p}(x, y)$ on a general reference surface. There may be points on the surface that do not have incoming rays. At other points, several rays may pass through the same point.

For many imaging systems, there is a distinguished ray surface which is defined *solely* by the geometry of the rays. In Figure 5, we see the envelope of the incoming rays form a curve. On such surfaces, direction \mathbf{q} is really a function of position \mathbf{p} , and the incoming rays are tangent to the surface. This is a special case of a ray surface called a *caustic*. We argue that caustics are the logical place to locate our raxels.

Caustics have a number of different characterizations. The caustic often forms a surface to which all incoming rays are tangent. This does not happen in the perspective case, where the caustic is a point. Caustics can also be viewed as geometrical entities where there occurs a singularity (bunching) of the light rays. As a result, a caustic formed by rays of illumination generates a very bright region when it intersects a surface.⁵ In our context of imaging systems, a caustic is the locus of points where incoming rays most converge. Points on the caustic can therefore be viewed as the ‘locus of view-points’ of the system. It is thus a natural place to locate the raxels of our abstract imaging model.

⁵For example, when light refracts through shallow water of a pool, bright curves can be seen where the caustics intersect the bottom.




Figure 5: The caustic is a good candidate for the ray surface of an imaging system as it is closely related to the geometry of the incoming rays; the incoming ray directions are tangent to the caustic.

3.1 Definition of a Caustic Surface

In section 2.3, we described the set of points and directions that can be detected by the imaging system. This set is described by the range of $\mathcal{L}(x, y, r)$ from equation (2). The position component functions are $X = p_X(x, y, r)$, $Y = p_Y(x, y, r)$, and $Z = p_Z(x, y, r)$. The map from (x, y, r) to (X, Y, Z) can be viewed as a change of coordinates. The *caustic surface* is defined as the locus of points in X, Y, Z space where this change of coordinates is singular.

3.2 Computing Caustics

The ray-to-image mapping of an imaging system may be obtained in two ways. One way, is to derive the mapping or compute it numerically from optical components or parameters of an imaging system which is known *a-priori*. Alternatively, a calibration method can be used to obtain a discrete form of the mapping (see section 5). In either case, our goal is to compute the caustic surface from the given mapping.

When this mapping is known in closed form, analytic methods can be used to derive the caustic surface [4]. When a discrete approximation is given, a host of numerical methods [12], [18], [26], may be used. The method we use computes the caustic by finding all the points where the change in coordinates described above is singular [2], [3].

Equation (2) expresses \mathcal{L} in terms of a known or measured ray surface $\mathbf{s}(x, y)$. The caustic is defined as the singularities in the change from (x, y, r) coordinates to (X, Y, Z) coordinates given by \mathbf{p} . Singularities arise at those points (X, Y, Z) where the Jacobian matrix J of the transformation does not have full rank. We can find these points by

computing the determinant of the Jacobian,

$$\det(J) = \begin{vmatrix} \frac{\partial p_X}{\partial x} + r \frac{\partial q_X}{\partial x} & \frac{\partial p_X}{\partial y} + r \frac{\partial q_X}{\partial y} & q_X \\ \frac{\partial p_Y}{\partial x} + r \frac{\partial q_Y}{\partial x} & \frac{\partial p_Y}{\partial y} + r \frac{\partial q_Y}{\partial y} & q_Y \\ \frac{\partial p_Z}{\partial x} + r \frac{\partial q_Z}{\partial x} & \frac{\partial p_Z}{\partial y} + r \frac{\partial q_Z}{\partial y} & q_Z \end{vmatrix}, \quad (3)$$

and setting it to zero. Since this is quadratic in r we can solve for r explicitly in terms of \mathbf{p} , \mathbf{q} , and their first derivatives with respect to x and y . Plugging this back into \mathcal{L} gives us an expression for the caustic ray surface parameterized by (x, y) as in equation (1). If the optical system has translational or rotational symmetry then we may only consider one parameter, for example x , in the image plane. In this case Jacobian becomes linear in r and the solution simplifies to:

$$r = \frac{(q_X \frac{dp_Y}{dx} - q_Y \frac{dp_X}{dx})}{(q_Y \frac{dq_X}{dx} - q_X \frac{dq_Y}{dx})}. \quad (4)$$

3.3 Field of View

Some parameters used to specify the general perspective camera model are derived from the ray surface representation in our general imaging model. Other parameters depend on the perspective assumption and are ill defined in our model. For example, field of view presents an ambiguity, since in the non-perspective case the rays may no longer form a simple rectangular cone. One candidate for a field of view is the range of $\mathbf{q}(x, y)$ over the image. This is the same as the Gauss map [11]. The Gauss map is a good approximation to the field of view when the scene points are distant relative to the size of the imaging system.

Other geometric parameters of a perspective imaging system, such as aspect ratio, and spherical aberration [25], may no longer be separable in the general imaging model from the parameterized ray surface.

4 Non-Geometric Raxel Parameters

Each raxel is treated as a very narrow field of view, perspective imaging system. Many of the conventional parameters associated with a perspective or near-perspective systems may be attributed to a raxel.

4.1 Local Focal Length and Point Spread

An arbitrary imaging system cannot be expected to have a single global focal length. However, each raxel may be modeled to have its own focal length. We can compute each raxel's focal length by measuring its point spread function for several depths. A flexible approach models the point spread as an elliptical Gaussian. Each ellipse has a major and minor axis. The major axis makes an angle ψ with the x -axis in the image.⁶ Each raxel has two focal lengths, f_a ,

⁶The angle ψ is only defined if the major and minor axis have different lengths.

f_b , as well as a focal orientation ψ in the image.

4.2 Radiometry

The radiometric response, g , is expected to be smooth and monotonic and can be modeled a polynomial [19]. If one can compute the radiometric response function of each raxel, one can linearize the response with respect to scene radiance, assuming the response is invertible.

We model the relation between raxel irradiance E as scene radiance L times a spatially varying attenuation factor, $h(x, y)$, corresponding to the image point (x, y) . We call $h(x, y)$ the *fall-off function*. This factor takes into account the finite size of the aperture and vignetting effects which are linear in L .

4.3 Complete Imaging Model

The general imaging model consists of a set of raxels parameterized by x and y in pixel coordinates. The parameters associated with these raxels (see Figure 6), are (a) position and direction, that describe the ray surface of the caustic, (b) major and minor focal lengths as well as a focal orientation, (c) a radiometric response function, and (d) a fall-off constant for each pixel.

x	y	p_X	p_Y	p_Z	q_θ	q_ϕ	f_a	f_b	Ψ	g	h
-----	-----	-------	-------	-------	------------	----------	-------	-------	--------	-----	-----

Figure 6: Each raxel has the above parameters of image coordinates, position and direction in space, major and minor focal lengths, focal orientation, radiometric response, and fall-off factor. These parameters are measured with respect to a coordinate frame fixed to the imaging system.

Camera parameters are separated into external and internal parameters. The external parameters are specified by a coordinate frame. The internal parameters (Figure 6) are specified with respect to that coordinate frame. In particular, for each raxel i , the parameters \mathbf{p}_i , \mathbf{q}_i are measured with respect to a single coordinate frame fixed to the imaging system. If the system is rotated or translated, these parameters will not change but the coordinate frame will.

In the case of perspective projection, the essential [5] or fundamental [10] matrix provides the relationship between points in one image and lines in another image (of the same scene). In the general imaging model, this correspondence need no longer be projective. Nevertheless, a point in one image still corresponds to a curve in the other, which may be computed.

5 Finding the Model Parameters

In section 3 we described how to compute our model for a known optical system. In contrast, our goal in this section




Figure 7: (a) A ray corresponding to a pixel i intersects two planes, separated by a known distance Δz , at points \mathbf{p}_n and \mathbf{p}_f . If these positions are known, the direction of the ray \mathbf{q}_f may be determined for each pixel. From this we compute the raxel parameters. To determine from the image the positions of \mathbf{p}_n and \mathbf{p}_f we place an active display at the planes and use binary coding patterns. By rotating the image device we may perform ray-based calibration for imaging systems with an arbitrary field of view. (b) One of the binary patterns used for ray-based calibration.

is to find the parameters for an unknown (black box) imaging system. Now, we construct a calibration environment where the geometric and radiometric parameters can be efficiently estimated. The geometric parameters are determined by establishing the image point to scene ray correspondence or *ray based calibration*.⁷ We can find the radiometric response function and the fall-off by controlling the intensity of light along the rays and measuring the response in the image.

Suppose we could determine the direction \mathbf{q}_f and position \mathbf{p}_f of each point on a plane imaged at i (Figure 7(a)). This determines a ray surface for the imaging system. From this we compute the caustic ray surface and thus the raxel parameters for each point in the image. We may determine the direction \mathbf{q}_f on the plane by finding a point \mathbf{p}_n , also imaged by i , at a second plane translated a known distance Δz toward the imaging device.

To determine the position on a plane imaged at i we place an active display at the plane. If a display has N locations, we can make each point distinct in $\log N$ images using simple grey coding or bit coding. A million pixel display requires only twenty images to encode every point. This two plane method solves the ray based calibration problem for a limited field of view or a rotationally symmetric system. We may calibrate the remaining imaging systems by placing them on a turntable and rotating the coordinate frame of the imaging system.

We display a linearly increasing constant brightness sequence to our imaging system. First, we calibrate the radiometric response function for a representative point in the

⁷It is important to note that, given the non-perspective nature of a general device, conventional calibration [5], [10], [15], [25], cannot be applied here.

image from the known input brightness. We may then compute the fall-off function across all the points.

We used twenty binary coded images for each plane for ray based calibration (assuming a one megapixel display). We compute both the radiometric response function and the fall-off from seventeen uniform brightness levels. Thus, with roughly 60 images we have measured the parameters of our general imaging model, with the exception of point spread. To make those 60 images robust to noise each image was an average of 30 images of the same display input.

5.1 Experimental Apparatus

We calibrated two imaging systems to demonstrate the generality of our method. One imaging system was a non-perspective catadioptric system consisting of a perspective camera and a parabolic mirror [23]. The second imaging system was a perspective camera. In both experiments the camera was a Cannon Optura digital video camera. Bit patterns were displayed on a laptop with a 14.5 inch LCD screen with resolution 1024×768 pixels. Rather than move the display, the camera was mounted on a stage which was translated 60mm in the direction normal to the screen.

In Figure 8(a) the parabolic catadioptric system is shown.⁸ The laptop was oriented so as to give the maximum screen resolution along the axis of symmetry. Figure 8(b) shows a sample binary pattern as seen from the parabolic catadioptric system. The perspective imaging system, consisting of just the camera itself, can be seen in Figure 11(a). Figure 11(b) shows an image of a pattern of vertical bars from the perspective camera.

5.2 Experimental Results: Geometric Parameters

Figure 9 shows the recovery of the caustic for the non-perspective catadioptric system using ray based calibration. Since the system is rotationally symmetric, the image to ray map was recovered along a radius. The caustic was computed using equation (4). Since there is no data along the axis of symmetry, its position was estimated. The caustic of the system should have a cross section similar to the curve shown in Figure 5 (see [23] for details). The recovered caustic matches that part of the curve near the cusp. We only see this part of the caustic because the outer field of view of the imaging system ends where the rays are nearly normal to the mirror's axis of symmetry. Near the axis of symmetry the radial partial derivative increases faster. This means that near the axis, smaller changes in position on the caustic yield larger changes in angle. Thus, there is a drop in resolution near the axis of symmetry.

We solved for the zeros of the three dimensional Jacobian of equation 3 to recover the caustic scatter plot as seen in

⁸The parabolic mirror had an outer diameter of 100mm and an inner diameter of 3mm. The focus of the parabola was 25mm from the base.

Figure 12. Since the caustic of a perspective camera is the center of projection, the point-like measured caustic agrees with expectation.

5.3 Experimental Results: Non-Geometric Parameters

Figure 10(a) shows the normalized radiometric response function. As is the case with higher quality digital imaging devices, the response is close to linear over much of its dynamic range. The curve shown is a polynomial fit with the non-linear constraint that the curve must have a positive derivative.

We linearized the response and calculated the radiometric fall-off for the non-perspective case (Figure 10 (b)). The fall-off is normalized so the maximum is unity. The directionality of the LCD irradiance has been measured and normalized to be uniform with respect to direction. The pronounced fall-off away from the axis of symmetry is due to the higher resolution of non-perspective system there. Image pixels away from the axis of symmetry see a smaller object area. Hence, they gather less light from a uniform radiator.

For the perspective system, we might expect a $\cos^4 \alpha$ type fall-off [11]. However, many lens systems have been designed to remove this effect. Indeed, we find no significant fall-off as the function is nearly constant. This can be seen in Figure 13(a) where intensity represents the fall-off function. The near constancy of the fall-off is clearer in the graph of a radial slice 13(b). To conclude, this simple calibration procedure allows us to compute the parameters of our general and flexible model.

6 Acknowledgments

This work was supported in parts by a National Science Foundation ITR Award, IIS-00-85864, and a DARPA/ONR MURI Grant, N00014-95-1-0601. The authors would also like to thank Rahul Swaminathan, Yoav Schechner, Srini-vasa Narasimhan, and Joshua Gluckman for their comments.

References

[1] E. H. Adelson and J. R. Bergen. *The Plenoptic Function and the Elements of Early Vision*. The MIT Press, 1991.
 [2] M. Born and E. Wolf. *Principles of Optics*. Pergamon Press, 1965.
 [3] D. G. Burkhard and D. L. Shealy. Flux density for ray propagation in geometrical optics. *Journal of the Optical Society of America*, 63(3):299–304, March 1973.
 [4] S. Cornbleet. *Microwave and geometrical optics*. Academic Press, London, 1994.

[5] O. D. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In *Proc. ECCV*, pages 563–578, 1992.
 [6] N. Franceschini, J. M. Pichon, and C. Blanes. From insect vision to robot vision. *Philosophical Transactions of the Royal Society London: Biological Sciences (Series B)*, 337:283–294, 1992.
 [7] E. Gaten. Geometrical optics of a galatheid compound eye. *Journal Comparative Physiology, A.*, 175:749–759, 1994.
 [8] A. Gershun. Svetovoe pole (the light field, in english). *Journal of Mathematics and Physics*, XVIII:51–151, 1939.
 [9] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The lumigraph. In *Computer Graphics, Proc. SIGGRAPH*, page 43, 1996.
 [10] R. Hartley. Estimation of relative camera positions for uncalibrated cameras. In *Proc. ECCV*, pages 579–587, 1993.
 [11] B. Horn. *Robot Vision*. The MIT Press, 1986.
 [12] H. W. Jensen. Rendering caustics on non-lambertian surfaces. *Computer Graphics Forum*, 16(1):57–64, March 1997.
 [13] M. J. Langer and S. W. Zucker. What is a light source? In *Proc. CVPR*, pages 172–178, 1997.
 [14] M. Levoy and P. Hanrahan. Light field rendering. In *Computer Graphics, Proc. SIGGRAPH*, page 31, 1996.
 [15] H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, September 1981.
 [16] S. Mann. Pencigraphy with AGC : Joint parameter estimation in both domain and range of functions in same orbit of the projective-wyckoff group. <http://citeseer.nj.nec.com/36900.html>, 96.
 [17] D. McCutchen. Method and Apparatus for Dodecahedral Imaging System. *U. S. Patent 5,023,725, June 11, 1991.*, 1991.
 [18] D. Mitchell and P. Hanrahan. Illumination from curved reflectors. *Computer Graphics*, 26(4):283–291, 1992.
 [19] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Proc CVPR*, volume 2, pages 374–380, June 1999.
 [20] K. Miyamoto. Fish eye lens. *Journal of Optical Society of America*, 54(8):1060–1061, August 1994.
 [21] S. K. Nayar and S. Baker. Catadioptric image formation. In *DARPA97*, pages 1431–1438, 1997.
 [22] S. Peleg, Y. Pritch, and M. Ben-Ezra. Cameras for stereo panoramic imaging. In *Proc. CVPR*, volume 2, pages 201–214. IEEE Computer Society, June 2000.
 [23] R. Swaminathan, M. D. Grossberg, and S. K. Nayar. Caustics of Catadioptric Cameras. In *these proceedings, Proc. ICCV*, 2001.
 [24] R. Swaminathan and S. K. Nayar. Nonmetric calibration of wide-angle lenses and polycameras. *PAMI*, 22(10):1172–1178, October 2000.
 [25] R. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision. *International Journal of Robotics and Automation*, 3(4):323–344, Aug 1987.
 [26] M. Watt. Light-water interaction using backward beam tracing. *Computer Graphics*, 24(4):377–385, 1990.




Figure 8: (a) A non-perspective imaging system and a calibration system, consisting of a laptop and a translating stage. The non-perspective catadioptric system consists of a perspective camera a parabolic mirror. The imaging system is mounted on the translating stage. The laptop displays 26 patterns. (b) A sample bit pattern as seen through the parabolic catadioptric system.




Figure 9: The caustic recovered for a parabolic catadioptric system. A cross section of the caustic of the system is similar to the central part of the curve in Figure 5.




Figure 10: (a) The normalized radiometric response is calculated from images of 17 uniform screens. The curve is a polynomial fit with the endpoints constrained and the first derivative required to be positive. (b) The radial fall-off function of the parabolic catadioptric system. The plot goes from the edge toward the center of the system.




Figure 11: (a) A perspective imaging system and a calibration system, consisting of a laptop and a translating stage. The axis of perspective camera is normal to the plane of the screen. (b) A sample bit pattern as seen through the perspective system.




Figure 12: The perspective caustic is a small cluster of points corresponding to the center of projection. The scale is approximately the same as Figure 9 for comparison.




Figure 13: (a) The perspective two dimensional fall-off function as an image. (b) A plot of the radial fall-off function for the perspective system. The plot goes from the edge to the center of the image since it is symmetric.