

Spherical Catadioptric Arrays: Construction, Multi-View Geometry, and Calibration

Douglas Lanman* Daniel Crispell* Megan Wachs[†] Gabriel Taubin*

*Brown University, Providence, Rhode Island, USA

[†]Stanford University, Stanford, California, USA

Abstract

This paper introduces a novel imaging system composed of an array of spherical mirrors and a single high-resolution digital camera. We describe the mechanical design and construction of a prototype, analyze the geometry of image formation, present a tailored calibration algorithm, and discuss the effect that design decisions had on the calibration routine. This system is presented as a unique platform for the development of efficient multi-view imaging algorithms which exploit the combined properties of camera arrays and non-central projection catadioptric systems. Initial target applications include data acquisition for image-based rendering and 3D scene reconstruction. The main advantages of the proposed system include: a relatively simple calibration procedure, a wide field of view, and a single imaging sensor which eliminates the need for color calibration and guarantees time synchronization.

1. Introduction

Catadioptric systems are optical systems composed of both mirrors and lenses. Most catadioptric cameras used in computer vision applications contain only one or two curved mirrors. In this paper we study an optical system composed of a large number of identical spherical mirrors forming a regular planar tiling, imaged by a single high resolution sensor. We describe the construction of the prototype shown in Figure 1, we analyze the multi-view geometry of image formation, and we propose a calibration algorithm. The literature on catadioptric systems is extensive [5, 3, 4, 11, 13, 14, 17] but to our knowledge, systems with large numbers of identical mirrors arranged in regular configurations have not been presented. Other methods for simultaneously acquiring multiple images from different viewpoints have included arrays of cameras and moving cameras [9, 6, 2, 19, 18] as well as inexpensive arrays of lenses mounted on flatbed scanners [20]. The main advantages of our system are the wide field of view and single-image data capture, which makes the time synchro-

nization issues associated with multi-camera systems vanish. On the other hand, frame-rate video processing is currently not possible with the high resolution consumer-grade digital camera used in our prototype.

Most calibration algorithms for catadioptric systems are designed for single mirror systems [8, 21]. More recently, non-parametric general imaging models have been proposed to describe non-pinhole cameras, as well as for cal-



Figure 1: Spherical catadioptric array. (A) Mechanical design. (B) Constructed prototype. (C) Typical image captured by the 8 megapixel digital camera. The four corner pins marked in red are used to estimate the plate coordinate system (as discussed in Section 3).

*{douglas.lanman, daniel.crispell, taubin}@brown.edu

ibrating such systems [7, 16, 15]. It is appropriate to use these methods when the geometry of the optical system components is unknown. In systems such as the one presented in this paper, where complete geometric design data is available, a parametric description of the process of image formation produces better results.

The paper is organized as follows: in Section 2 we describe the construction of the system. Section 3 presents the calibration algorithm. In Section 4 we review the geometry of image formation for spherical catadioptric systems. Section 5 presents the multi-view geometry equations used for calibration and 3D reconstruction. In Section 6 we describe implementation details and calibration results. Section 7 discusses results for some preliminary applications. Finally, in Section 8, we elaborate on our future plans.

2. Mechanical Design and Construction

We designed and built the system shown in Figure 1. It consists of a black anodized aluminum plate, 1/4" thick, with cylindrical stainless steel pins pressed into holes. These holes, drilled at precise locations on the plate, are used to hold and align 31 spherical mirrors. The pins are cut to equal length and inserted into the plate with high precision, so that the pin heads are coplanar to machine shop precision. The inexpensive plastic mirrors shown in Figure 2 are glued to the aluminum plate using a synthetic silicon rubber adhesive. This mounting method prevents damage to the reflective layer, but it also causes inaccuracies in the positioning of mirrors with respect to the plate. As a result, the mirror parameters (i.e., the location of the sphere centers and their radii) are not known in advance with high precision, and must be refined by the calibration algorithm. The plate is positioned in space to roughly fill the field of view of an Olympus C-8080 8 megapixel digital camera and at an angle such that the camera cannot see its own image in any of the mirrors. The structure shown in Figure 1(A), built out of standard industrial aluminum extrusions, is used as the system skeleton. A single image captures all 31 mirrors. Figure 1(C) shows a typical image captured by the camera. The Olympus SDK provided by the manufacturer [12] is used to automate the capture and calibration processes. Captured images are downloaded directly to the computer through the USB 2.0 interface.

3. Calibration

We use a pinhole perspective camera model for the digital camera and select its local coordinate system as the world coordinate system. The purpose of the calibration procedure is to determine the precise mapping from camera pixel coordinates to light rays in three-dimensional space [7, 16]. We divide the calibration process into three steps:



Figure 2: Spherical mirrors. (A) Adhesive-backed rear view mirrors used in the construction of the mirror array. (B) Detail of mirror array plate showing the steel pins that hold the mirrors in place.

(1) intrinsic calibration of the digital camera, (2) estimation of the plate pose with respect to the world coordinate system, and (3) calibration of the spherical mirrors with respect to the plate. The last step includes estimating the location of the sphere centers and their radii, as well as the pose of a calibration object. In a more complex model, it also includes estimating the thickness of the refractive layer in the mirrors, as well as their index of refraction of the protective mirror coating.

3.1 Intrinsic pinhole camera calibration

For this step we use the well-established Camera Calibration Toolbox for Matlab [1]. This step also includes distortion correction for the camera lens.

3.2 Plate pose estimation

This step is completed using a single input image and does not require a calibration object. In fact, the reflected images in each mirror are ignored. An ellipse detection algorithm is used to locate the four pins closest to the four corners of the image, marked in red in Figure 1(C). Conservative search areas for these pins can be predicted from the mechanical design data. Since we know the precise location of these four pins in the plate and their relative distances, we compute a homography from the four point correspondences, and from this homography we obtain a first estimate of the equation of the plane in space. We subsequently predict the locations of the rest of the pins in the image. These predictions are then used to search for the remaining pins. The same ellipse detection algorithm is used to refine the location estimates for all the remaining pins; afterwards, we recompute a more precise homography from this data. This step has to be repeated every time the camera moves with respect to the plate, or when intrinsic camera parameters (e.g., focal length) are changed.

3.3 Calibration of spherical mirrors

This step requires a single image of a calibration target. We use a standard planar checkerboard pattern which we position so that most of its corners are visible in the reflection from every mirror surface. We model the mirrors as perfect spheres with centers lying approximately in a plane located behind the plane of the steel pin heads. In this step, we determine the location of the spherical mirror centers (with respect to the plate and with respect to the world coordinate system), the radii of the mirrors, and the location of the calibration target in world coordinates. Initial estimates for the sphere centers and sphere radii are available from the mechanical design data, but the exact radii and precise center locations are not known because of the imprecise mounting method used.

3.3.1 Bundle adjustment

We propose a straightforward method for the final calibration step based on a new bundle adjustment formulation. In Section 4 we derive the equations of image formation for a catadioptric system composed of a pinhole camera and a spherical mirror. These equations determine a mapping which assigns an image point u to a directed *reflected ray*

$$R_u = \{p = q(u, \Lambda_i) + \lambda v(u, \Lambda_i) : \lambda > 0\} \quad (1)$$

incident on the i th spherical mirror, where Λ_i is a set of parameters including the sphere center and radius, and p is a 3D point whose image on the i th mirror is u , and $\lambda > 0$ is a constant. A similar parametric form, but with different parameter values Λ_j , describes the process of image formation for the j th mirror.

For every identifiable point p_j in a calibration target (e.g., a checkerboard corner) and each corresponding image point u_{ij} associated with the i th mirror, we have one ray equation

$$q(u_{ij}, \Lambda_i) + \lambda_{ij} v(u_{ij}, \Lambda_j) = q_{ij} + \lambda_{ij} v_{ij} = p_j. \quad (2)$$

All of these equations must be satisfied simultaneously, but due to measurement errors they can only be satisfied in the least squares sense. In Section 5 we show that this problem is equivalent to minimizing the error function

$$E(P, \Lambda) = \sum_{(i,j) \in \mathcal{I}} \|v(u_{ij}, \Lambda_i) \times (p_j - q(u_{ij}, \Lambda_i))\|^2 \quad (3)$$

where $P = \{p_1, \dots, p_N\}$, the unknowns λ_{ij} have been eliminated, and \mathcal{I} is the set of pairs (i, j) such that point p_j is visible in the i th mirror.

3.3.2 Estimating point locations

In general, minimizing $E(P, \Lambda)$ is a non-linear least squares problem which requires initial estimates for both the point locations and the intrinsic parameters. Note that the mechanical design data provides accurate initial estimates for the intrinsic parameters, however the point location estimates are not available. We compute initial estimates for the point positions using the following procedure, which uses the initial estimates for the intrinsic parameters.

We show in Section 5 that if the intrinsic mirror parameters Λ are known, the location of the point p_j can be estimated in homogeneous coordinates as the right singular vector associated with the smallest singular value of a matrix W_{p_j} . We use the SVD algorithm to compute initial estimates for the points using the initial estimates for the intrinsic parameters. If the point p_j is visible in all the mirrors, then the matrix W_{p_j} has the following expression

$$W_{p_j} = \begin{pmatrix} \hat{v}_{1j} & -\hat{v}_{1j}q_{1j} \\ \vdots & \vdots \\ \hat{v}_{Nj} & -\hat{v}_{Nj}q_{Nj} \end{pmatrix}.$$

If not, then the rows corresponding to mirrors where p_j is not observed should be removed.

3.3.3 Rigid motion model

A local descent algorithm can be used to refine the point locations and intrinsic spherical mirror parameter estimates. In our experience, this procedure does not produce accurate results, mainly because some points are only visible in a few mirrors. The main limitation of this approach is that it does not impose as a constraint the fact that the spatial relation between the points p_j is known.

Since calibration features p_j are points on a known calibration object, which in our case are the corners of a planar checkerboard, we can write each point with respect to a coordinate system that describes the pose of the calibration object

$$p_j = R \xi_j + T,$$

where R is a three-dimensional rotation matrix, T a three-dimensional translation vector, and ξ_j is a known three-dimensional vector which describes the coordinates of the point p_j in the calibration object coordinate system. As usual, for a planar object such as a checkerboard pattern, we can choose the local coordinate system such that one coordinate of the vectors ξ_j is equal to zero. With this substitution, the $3N$ free variables in the energy function corresponding to the coordinates of the points p_j are reduced to just 6, resulting in a new energy function $E(R, T, \Lambda)$.

Again, to locally minimize this new energy function, we need initial estimates. We already have initial estimates for

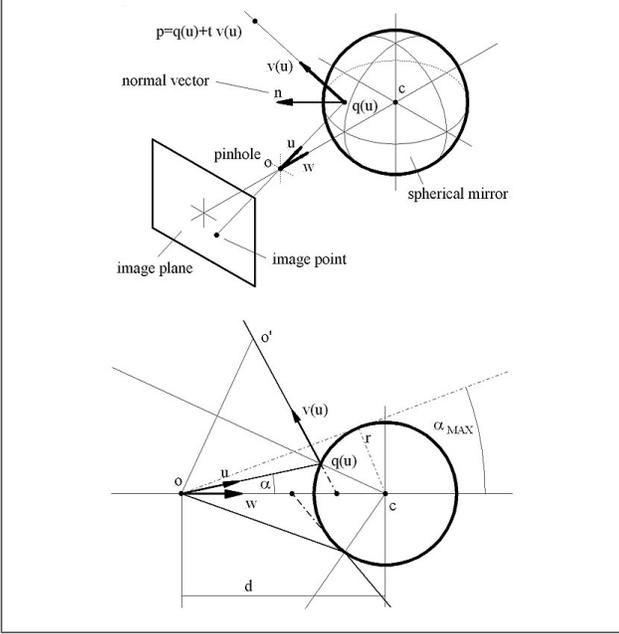


Figure 3: The geometry of a reflected ray. The incident ray, normal vector, and reflected ray are coplanar.

Λ . These are the same ones used in the previous formulation. To compute initial estimates for R and T , we choose three points p_j, p_k, p_l which are seen in most mirrors, and which are not very close together on the checkerboard. We use the SVD-based estimation algorithm described above to estimate their positions. From these three positions we can estimate the pose (R, T) . Then we locally minimize the energy function $E(R, T, \Lambda)$ with respect to all the variables simultaneously.

Since the mirrors are not expected to move with respect to the plate, this process needs to be performed only once in a laboratory environment.

4 The Ray Equations

In this Section we derive an analytic form for the mapping from pixel coordinates to 3D reflected rays with respect to the world coordinate system. We first analyze the case of a first surface mirror (i.e., a mirror without a layer of transparent material in front of the reflective coating). Afterwards, we modify this model to account for refraction.

4.1 Model for a first surface mirror

The reflection of rays by a spherical mirror is governed by the classical law of reflection: the incident ray, the normal vector to the mirror surface at the point of incidence, and the reflected ray are coplanar. In addition, the incident and reflected rays are symmetric with respect to the normal

vector. This means that the optical system has cylindrical symmetry with respect to the line joining the pinhole and the center of the spherical mirror. We call this line the *optical axis* of the system. This symmetry is broken if the image plane is not orthogonal to this line. This is not a problem because the images formed on two different image planes are related by a homography. In fact, what is important is not where the ray hits an image plane, but its direction as it passes through the pinhole. We avoid this uncertainty by representing image points as unit length vectors u , after correcting for intrinsic camera calibration. This is equivalent to saying that we use a pinhole camera with a spherical retina.

To analyze the process of image formation in this system we consider rays traveling in the opposite direction, from the camera through the pinhole to space, hitting the mirror, and then reflected off of the spherical mirror, as illustrated in Figure 3. Here we denote o the pinhole, c the center of the spherical mirror, d the distance between the pinhole and the center of the sphere, r the radius of the spherical mirror, and w the vector from o to c normalized to unit length, so that $c - o = dw$. The origin of the world coordinate system is o .

Let u be an arbitrary unit length vector, let α be the angle formed by the vectors u and w , and let I_u be the *incident ray* that goes through the pinhole in the direction of u .

Again, because of the law of reflection, the incident ray, the reflected ray, the normal vector, and the optical axis are all coplanar. In particular, the line supporting the *reflected ray* R_u intersects the optical axis. But not all these rays intersect the axis at the same point. This is illustrated in Figure 3. In fact, the location along the optical axis is only a function of the angle between the vectors u and w . Two vectors u_1 and u_2 are congruent modulo a rotation around the optical axis, if and only if their corresponding reflected rays intersect the optical axis at the same point. If u_1 and u_2 are not congruent modulo a rotation around the optical axis, the corresponding reflected rays are *twisted* (i.e., the two supporting lines are not coplanar).

The incident ray I_u hits the mirror surface at the point $q = q(u)$, and reflects off according to the law of reflection along a unit length vector $v = v(u)$. The reflected ray R_u can be written in parametric form as shown in Equation 1 above, where $v = v(u)$ is its direction. We can write $q = \xi u$, for some positive number ξ . Since $\|q - c\|^2 = r^2$, and $c - o = dw$, ξ must be the smallest positive real root of the quadratic equation

$$0 = \|\xi u - dw\|^2 - r^2.$$

Since and $\|u\|^2 = \|w\|^2 = 1$, we have

$$0 = \|\xi u - dw\|^2 - r^2 = \xi^2 - 2\xi d(u^t w) + d^2 - r^2.$$

Rearranging terms and completing squares

$$0 = (\xi - d(u^t w))^2 - (r^2 - d^2(1 - (u^t w)^2)). \quad (4)$$

Note that, if α is the angle formed by the vectors u and w , the maximum angle α_{MAX} is achieved when the incident and reflected rays are tangent to the sphere, in which case

$$0 < 1 - (u^t w)^2 = \sin(\alpha)^2 \leq \sin(\alpha_{\text{MAX}})^2 = r^2/d^2 ,$$

and Equation 4 has two real roots

$$\xi = d(u^t w) \pm \sqrt{r^2 - d^2(1 - (u^t w)^2)} .$$

Since $u^t w > 0$, the root with the + sign is clearly positive, but so is the one with the - sign, because

$$\begin{aligned} 0 < r < d &\Rightarrow \\ (d(u^t w))^2 - (r^2 - d^2(1 - (u^t w)^2)) &= d^2 - r^2 > 0 \Rightarrow \\ \xi = d(u^t w) - \sqrt{r^2 - d^2(1 - (u^t w)^2)} &> 0 . \end{aligned}$$

We conclude that

$$q(u) = \left(d(u^t w) - \sqrt{r^2 - d^2(1 - (u^t w)^2)} \right) u . \quad (5)$$

We define the vector v as the mirror image of the incident vector u with respect to tangent plane to the mirror surface at q . Nayar [11] first proposed this derivation. To obtain an expression for v we decompose the vector u as the sum of a normal vector to the mirror at the point q , and an orthogonal vector

$$u = nn^t u + (I - nn^t) u ,$$

and we invert the normal component

$$v = -nn^t u + (I - nn^t) u = (I - 2nn^t) u .$$

4.2 Model for a second surface mirror

The inclusion of a refractive layer in front of the mirror surface only requires a slight modification to the model outlined above. Surprisingly, a second surface spherical mirror can be modeled as a first surface mirror with an identical center and a modified radius ($r' = r + \delta$). This result is graphically demonstrated in Figure 4. Here, u once again denotes the direction of the ray traveling from the camera pinhole o to the mirror.

At the first mirror surface, we must apply Snell's Law

$$\sin(\alpha) = \nu \sin(\beta) , \quad (6)$$

where α is the angle of incidence (with respect to the surface normal n), β is the angle of refraction, and ν is the index of refraction. Given the refracted ray, we can solve for its intersection with the second mirror surface, located a distance r from the mirror center. As demonstrated in Figure 4, the path of the ray after reflection at the second surface must be symmetric about the line s connecting the mirror center to the point of intersection with the second surface. This symmetry property is exploited to complete the ray diagram.

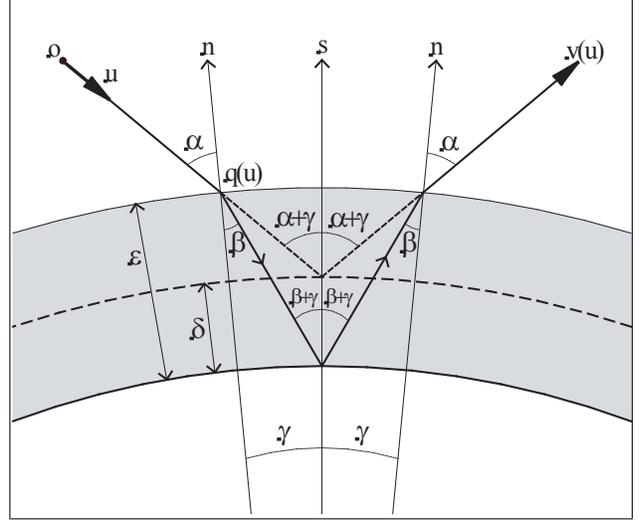


Figure 4: Refraction diagram for a second surface spherical mirror. The shaded region denotes the refractive layer between the first and second surfaces.

Note that if we ignored the refraction at the first surface, the incident ray u would intersect the symmetry axis at $r' = r + \delta$. If there was a first surface mirror with radius r' centered at c , it would reflect this ray such that it would coincide with the outgoing ray from the second surface mirror. As a result, if we can solve for δ , then we can model second surface mirrors by their equivalent first surface representations.

In general, δ is a function of the angle of incidence α and the properties of the refractive layer: the index ν and the thickness ϵ . Inspecting Figure 4, we find that the law of sines produces the following equalities

$$(1 + \epsilon) \sin(\alpha) = (1 + \delta) \sin(\alpha + \gamma) \quad (7)$$

$$(1 + \epsilon) \sin(\beta) = (1 + \delta) \sin(\beta + \gamma) , \quad (8)$$

where γ is the angle between the refracted ray and the symmetry axis s .

Using Equations 6, 7, and 8, we can solve for $\delta(\alpha, \epsilon, \nu)$ by eliminating the dependence on $\{\beta, \gamma\}$. From 7, we have

$$\delta = \frac{(1 + \epsilon) \sin(\alpha)}{\sin(\alpha + \gamma)} - 1 . \quad (9)$$

The denominator $\sin(\alpha + \gamma)$ can be written as a function of $\{\alpha, \epsilon, \nu\}$ as follows. First, note that $\sin(\alpha + \gamma) = \sin[(\beta + \gamma) + (\alpha - \beta)]$. This yields four trigonometric functions with dependence on $\{\alpha, \beta, \gamma\}$

$$\begin{aligned} \sin(\alpha + \gamma) &= \sin(\alpha - \beta) \cos(\beta + \gamma) + \\ &= \cos(\alpha - \beta) \sin(\beta + \gamma) . \end{aligned}$$

Using the trigonometric identity, $\cos(\theta) = \sqrt{1 - \sin^2(\theta)}$, it is apparent that we only require a solution for $\sin(\alpha - \beta)$

and $\sin(\beta+\gamma)$. We can solve for $\sin(\beta+\gamma)$ using Equations 6 and 8, giving

$$\sin(\beta + \gamma) = \left(\frac{1 + \epsilon}{\nu} \right) \sin(\alpha) .$$

Similarly, application of Snell's law provides the following solution for the remaining term

$$\sin(\alpha - \beta) = \frac{\sin(\alpha)}{\nu} \left[\sqrt{\nu^2 - \sin^2(\alpha)} - \sqrt{1 - \sin^2(\alpha)} \right] .$$

Substituting these results into Equation 9, we obtain the desired solution for δ as a function of the known parameters $\{\alpha, \epsilon, \nu\}$.

In conclusion, we find that the model for a second surface mirror reduces to that of a first surface mirror, with a modified radius. In general, the "effective" radius ($r' = r + \delta$) of the second surface mirror will be dependent on the angle of incidence and can be calculated following the derivation given above.

5 Multi-Mirror View Geometry

Here we assume that we have a pinhole camera with a spherical retina, that the intrinsic calibration of the pinhole camera is known, and that N mirrors are visible within the camera field of view. Each spherical mirror has its own parameter vector Λ_i (containing the location of its center and radius, and optionally the refraction layer thickness and the index of refraction), and its own parametric ray equation $p = q(u, \Lambda_i) + t v(u, \Lambda_i)$.

Suppose that a three-dimensional point p is visible in all the mirrors. Let u_i be the pixel coordinates of the image of p produced by the i th mirror, and let's denote $q_i = q(u_i, \Lambda_i)$ and $v_i = v(u_i, \Lambda_i)$. The q_i 's and v_i 's are not independent. The following constraints, which generalize the epipolar constraints, must be satisfied for them to define rays intersecting at a common point:

$$\text{There exist } t_1, \dots, t_n \in \mathbb{R} \text{ and } p \in \mathbb{R}^3 \text{ such that } p = q_i + t_i v_i \text{ for } i = 1, \dots, N . \quad (10)$$

This is a system of linear equations which can be written in matrix form as follows

$$\begin{pmatrix} v_1 & \cdots & 0 & \cdots & 0 & -I & q_1 \\ \vdots & \ddots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & \cdots & v_i & \cdots & 0 & -I & q_i \\ \vdots & \ddots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & \cdots & v_N & -I & q_N \end{pmatrix} \begin{pmatrix} t_1 \\ \vdots \\ t_N \\ p \\ 1 \end{pmatrix} = 0 \quad (11)$$

where I is the three-dimensional identity matrix. Let's denote by N_p the left hand side $3N \times (N + 4)$ matrix. All

the elements of this matrix are functions of measurements and calibration parameters. If the calibration parameters are known, the position of the point p in space can be determined by solving this system of linear equations. Since $3N \geq N + 4$ for $N \geq 2$, in order for this linear system to have a non-trivial solution, the columns of N_p must be linearly dependent. Equivalently, the following rank inequality must hold

$$\text{rank}(N_p) \leq N + 3 ,$$

with equality for a unique solution. In practice, the matrix N_p is usually full-rank due to measurement errors and noise, and the problem must be solved in the least squares sense. To estimate p we can compute the right singular vector of N_p associated with its minimum singular value, and then normalize it so that the last coordinate is equal to unity. The three coordinates preceding the last one in this vector are then the coordinates of p .

A simpler approach is to first eliminate the parameters t_1, \dots, t_N by observing that an equivalent condition for all the rays to intersect at p is that

$$v_i \times (p - q_i) = 0 \text{ for } i = 1, \dots, N ,$$

which can be written in matrix form as

$$\begin{pmatrix} \hat{v}_1 & -\hat{v}_1 q_1 \\ \vdots & \vdots \\ \hat{v}_N & -\hat{v}_N q_N \end{pmatrix} \begin{pmatrix} p \\ 1 \end{pmatrix} = 0 \quad (12)$$

where \hat{v} is the 3×3 skew-symmetric matrix representing the vector product: $\hat{v} x = v \times x$, also called the covector [10]. Let's denote by W_p the left hand side $3N \times 4$ matrix of Equation 12. Again, the elements of the matrix W_p are functions of the measurements and calibration parameters. For this homogeneous linear system to have a non-trivial solution it is necessary that the matrix W_p satisfy the following condition

$$\text{rank}(W_p) \leq 3 ,$$

with equality for a unique solution. In practice, to estimate p we solve a linear least squares problem: we minimize the following expression with respect to p

$$E(p, \Lambda) = \sum_{i=1}^N \|v(u_i, \Lambda_i) \times (p - q(u_i, \Lambda_i))\|^2 \quad (13)$$

where $\Lambda = (\Lambda_1, \dots, \Lambda_N)$, by computing the right singular vector of W_p associated with the smallest singular value, and normalizing the solution so that the last coordinate is equal to unity.

6 Implementation and Results

The calibration process outlined in Sections 3 and 5 was implemented using Matlab. As previously discussed, our calibration procedure involves estimating a minimum of 100 parameters: the checkerboard position and orientation, the centers of all 31 mirrors, and a single mirror radius. In its most complete form, our model also accounts for the index of refraction, thickness of the dielectric layer, and independent mirror radii – leading to a maximum of 132 parameters. These parameters are estimated using Matlab’s `lsqnonlin` function to minimize Equation 13. In the following discussion, we will focus on the 100 parameter model, since it produced results nearly identical to the complete model.

We seed the iterative optimization process using the mirror positions estimated from the plate pose, as discussed in Section 3, and a coarse estimate of the mirror radius. Given these initial values, we determine the initial position of the calibration grid by solving Equation 11. Figure 5 shows the initial estimates obtained in a single calibration trial. Note that the initial mirror centers, as estimated from the plate pose, are biased from their optimized positions. This is due to an error in the initial guess for the mirror radius. This error also results, through the application of Equation 11, in a biased estimate of the checkerboard corners. These biases, however, are removed if a more accurate initial estimate of the mirror radius is available.

In Figure 6 we plot the optimized positions of the mirrors and the checkerboard. In addition, we use Equation 11 to once again reconstruct the individual checkerboard corners. Using the calibrated mirror parameters, the reconstructed checkerboard corners have a mean error of 2.3 mm when compared with the centrally-optimized estimates.

7 Applications

A spherical catadioptric array is ideally-suited for data acquisition in image-based rendering and 3D scene reconstruction. In these contexts our system possesses several desirable properties, including a relatively simple calibration procedure (outlined in this paper), a wide field of view, and a single imaging sensor – eliminating the need for color calibration and synchronization procedures. In order to motivate these applications, we will demonstrate the application of our system to 3D scene reconstruction.

As shown in Figure 7, a single image of a model house was acquired using the spherical mirror array. A simple mesh was constructed containing 11 vertices and 9 faces. Each vertex was manually labeled in each mirror image (if it was visible). Equation 11 was then applied to reconstruct the 3D position of all vertices. In order to assign color to each face, the ray equations, from Section 4, were used to

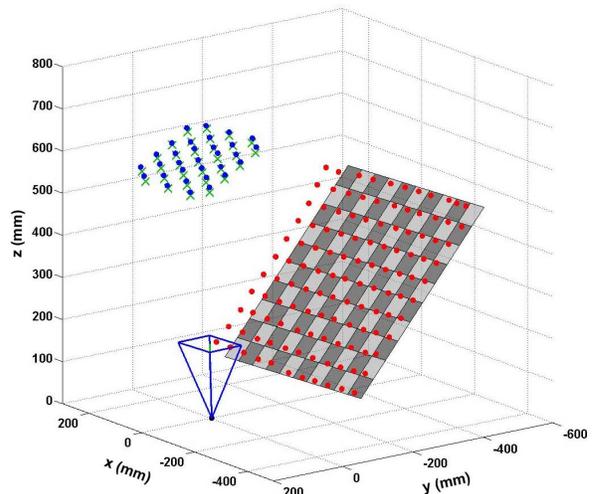


Figure 5: Calibration results. Red circles (●) represent initial estimates of checkerboard centers, green crosses (×) represent initial mirror centers, and blue circles (●) represent optimized mirror centers. The checkerboard is in its optimized position and orientation.

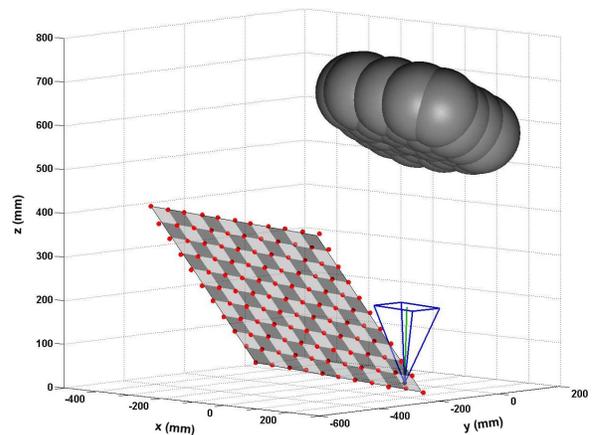


Figure 6: SVD-based reconstruction of the checkerboard. Red circles represent the reconstructed positions of the checkerboard corners. Individual mirrors are represented by gray spheres.

project pixels in each mirror image onto the reconstructed mesh. The resulting point cloud was rendering as an array of oriented, colored disks.

Although relatively simple, this example illustrates the inherent benefits of the proposed acquisition system, including automatic color calibration and frame synchronization. While pixels from five different mirrors were used to assign colors to the mesh faces, no significant color discrepancies were visually apparent in the synthesized model.

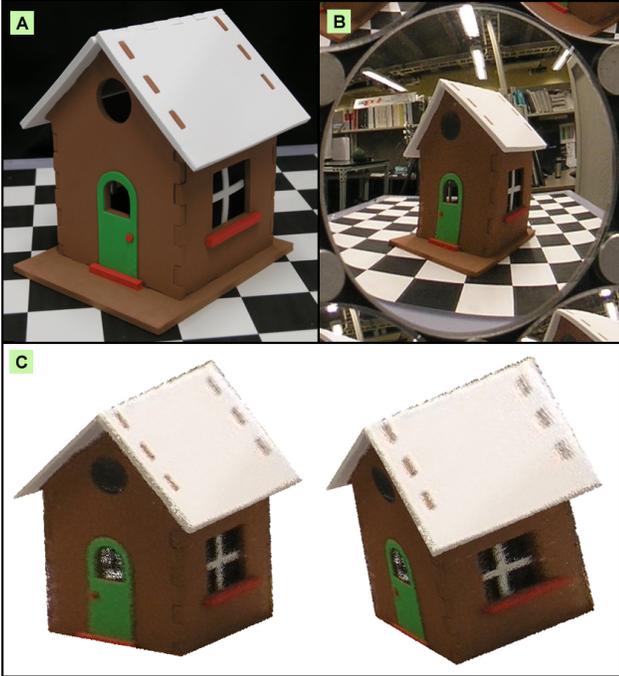


Figure 7: 3D reconstruction using a spherical catadioptric array. (A) Original scene. (B) Image of scene in spherical mirror. (C) 3D reconstruction results.

8 Future Work

In this paper we have discussed the general construction and application of an imaging system composed of an array of spherical mirrors and a single perspective camera. As discussed, this system was designed as a platform to develop more efficient multi-view imaging algorithms exploiting the combined properties of camera arrays and non-central projection catadioptric systems.

While preliminary experimental trials and reconstruction attempts proved promising, it is apparent that future studies will be hindered by the level of calibration precision currently provided by this system. One solution could involve using a hemispherical mirror, which would simplify calibration for single-mirror cases by allowing us to estimate all extrinsic parameters by simply determining the position of the silhouette of a sphere in an image. In addition, certain optical systems (e.g. high barrel distortion lenses) result in non-central imaging and may prove easier to calibrate.

In summary, we believe that inexpensive catadioptric arrays, such as the one described in this paper, are a viable acquisition platform for both the study and implementation of multi-view and non-central imaging algorithms. Future research should focus on refining both the construction and calibration of these novel acquisition platforms.

References

- [1] J.-Y. Bouguet. Camera Calibration Toolbox for Matlab. <http://www.vision.caltech.edu/bouguetj/calib.doc>.
- [2] C. Buehler, M. Bosse, S. Gortler, M. Cohen, and L. McMillan. Unstructured Lumigraph Rendering. In *Siggraph 2001, Conference Proceedings*, pages 425–432, 2001.
- [3] C. Geyer and K. Daniilidis. Catadioptric Camera Calibration. In *Proceedings of the 7th International Conference on Computer Vision*, 1999.
- [4] C. Geyer and K. Daniilidis. Catadioptric projective geometry. *International Journal of Computer Vision*, 45(3):223–243, 2001.
- [5] C. Geyer and K. Daniilidis. Properties of the Catadioptric Fundamental Matrix. In *Proceedings of the 7th European Conference on Computer Vision*, 2002.
- [6] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The Lumigraph. In *Siggraph 1996, Conference Proceedings*, pages 43–54, 1996.
- [7] M. Grossberg and S. Nayar. A General Imaging Model And A Method For Finding Its Parameters. In *ICCV*, 2001.
- [8] S. Kang. Catadioptric Self-Calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2000.
- [9] M. Levoy and P. Hanrahan. Light Field Rendering. In *Siggraph 1996, Conference Proceedings*, pages 31–42, 1996.
- [10] Y. Ma, S. Soatto, J. Kosecká, and S. S. Sastry. *An Invitation to 3-D Vision: From Images to Geometric Modeling*. Springer-Verlag, 2004.
- [11] S. Nayar. Sphero: Determining Depth using Two Specular Spheres and a Single Camera. In *Proceedings of SPIE Conference on Optics, Illumination, and Image Sensing for Machine Vision III*, pages 245–254, Nov. 1988.
- [12] Olympus SDK. <http://developer.olympus.com>.
- [13] A. Paulino and H. Araujo. Pose estimation for central catadioptric systems: an analytical approach. In *Proceedings of the 16th International Conference on Pattern Recognition*, pages 696–699, 2002.
- [14] R. Pless. Using Many Cameras as One. In *CVPR*, 2003.
- [15] P. Ramalingam, S. Sturm and S. Lodha. Towards complete generic camera calibration. In *CVPR*, 2005.
- [16] P. Sturm and S. Ramalingam. A Generic Concept For Camera Calibration. In *ECCV*, 2004.
- [17] R. Swaminathan, M. Grossberg, and S. Nayar. Caustics of Catadioptric Cameras. In *Proceedings of the 8th International Conference on Computer Vision*, 2001.
- [18] B. Wilburn, N. Joshi, V. Vaish, M. Levoy, and M. Horowitz. High speed videography using a dense camera array. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 294–301, 2004.
- [19] B. Wilburn, M. Smulski, H.-H. Kellin Lee, and M. Horowitz. The light field video camera. In *SPIE Electronic Imaging 2002, Media Processors, Conference Proceedings*, 2002.
- [20] J. Yang. Light fields on the cheap. Technical report, MIT, 2000.
- [21] X. Ying and Z. Hu. Catadioptric Camera Calibration Using Geometric Invariants. *IEEE Transactions On Pattern Analysis and Machine Intelligence*, 26(10), Oct. 2004.