Beyond Alhazen's Problem: Analytical Projection Model for Non-Central Catadioptric Cameras with Quadric Mirrors

Amit Agrawal, Yuichi Taguchi, and Srikumar Ramalingam
Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA
{agrawal, taguchi, ramalingam} at merl.com

Abstract

Catadioptric cameras are widely used to increase the field of view using mirrors. Central catadioptric systems having an effective single viewpoint are easy to model and use, but severely constraint the camera positioning with respect to the mirror. On the other hand, non-central catadioptric systems allow greater flexibility in camera placement, but are often approximated using central or linear models due to the lack of an exact model. We bridge this gap and describe an exact projection model for non-central catadioptric systems. We derive an analytical 'forward projection' equation for the projection of a 3D point reflected by a quadric mirror on the imaging plane of a perspective camera, with no restrictions on the camera placement, and show that it is an 8^{th} degree equation in a single unknown.

While previous non-central catadioptric cameras primarily use an axial configuration where the camera is placed on the axis of a rotationally symmetric mirror, we allow off-axis (any) camera placement. Using this analytical model, a non-central catadioptric camera can be used for sparse as well as dense 3D reconstruction similar to perspective cameras, using well-known algorithms such as bundle adjustment and plane sweeping. Our paper is the first to show such results for off-axis placement of camera with multiple quadric mirrors. Simulation and real results using parabolic mirrors and an off-axis perspective camera are demonstrated.

1. Introduction

A catadioptric sensor combines a camera with a mirror to increase the field of view (FOV) in a single photo. Such sensors typically use quadric shaped mirrors such as spherical, parabolic, elliptical, and hyperbolic, which are rotationally symmetric along an axis. However, catadioptric cameras are difficult to model due to the non-linear mapping of rays from the scene to the camera pixels. This non-linear mapping in general does not result in a single viewpoint (central) system and is difficult to model analytically.

Baker and Nayar [3] described the configurations for achieving a central catadioptric system (CCS), in which all the captured rays pass through a single virtual viewpoint inside the mirror. These include (a) an orthographic camera placed on the axis of a parabolic mirror, and (b) a perspective camera placed on the foci of a hyperbolic/elliptical mirror. Although a CCS is easy to model due to an effective single viewpoint, it severely restricts the placement of the camera with respect to the mirror. In addition, other configurations such as using a spherical mirror [20, 21, 14, 12, 26], using multiple mirrors [17, 8, 27] or off-axis placement of camera with respect to a quadric mirror [5] will lead to a non-central catadioptric system (NCCS). Finding a model for NCCS will significantly enhance the class of mirrors that can be used as well as lead to a simpler, easier and flexible hardware implementation. It will allow flexible mirror design to optimize for FOV, resolution characteristics or image-to-world mappings. In addition, such a model will allow fast sparse and dense 3D reconstruction using catadioptric cameras similar to perspective cameras.

An exact projection model is described by the problem of 'Forward Projection': What is the image projection of a 3D point reflected via a mirror? Previous approaches have acknowledged the difficulty of coming up with an analytical projection model and have resorted to optimization or search based techniques [20, 11, 19]. We refer to these techniques as iterative forward projection (IFP) [1]. In this paper, we solve this fundamental problem by deriving an analytical equation that describes the optical path from a 3D point to a perspective camera via reflection through a quadric mirror. We show that the projection of a 3D point for a NCCS with a quadric mirror, with no restriction on the camera placement, can be obtained by solving an 8^{th} degree polynomial equation in one unknown. Using this projection model, we show how sparse and dense 3D reconstruction can be performed for general catadioptric systems, similar to perspective cameras. In particular, we show 3D reconstruction using multiple parabolic mirrors and an off-axis perspective camera. Compared to IFP, we obtain a speed-up of 40X using our projection model.

Contributions: Our paper makes the following contributions:

- We describe an analytical solution of forward projection problem for NCCS with a quadric mirror, with no restriction on the camera placement.
- We demonstrate fast sparse and dense 3D reconstruction for NCCS with off-axis camera placement. Our approach avoids central or linear approximation and offers a speed up of $\sim 40X$ compared to previous iterative/search based techniques.

1.1. Related Work

Forward Projection (FP): The problem of determining the mirror reflection point for a spherical mirror is popularly known as the Alhazen's problem, first formulated by Ptolemy around 150 A.D. and later discussed extensively by Alhazen around 1000 A.D. Since then, this problem has fascinated mathematicians and several trigonometric as well as algebraic solutions have been proposed [2, 10]. While a spherical mirror has four solutions, it is assumed that there is no closed-form or analytical solution for general mirrors. Solving this problem analytically will be useful in several fields beyond catadioptric imaging, including rendering for computer graphics [6], image-based relighting [28], environment matting and shape from specular flow [23].

While analyzing defocus in catadioptric systems, Baker and Nayar [3] were unable to find an analytical solution for off-axis case and resorted to numerical methods. They require off-axis analysis since estimating the defocus blur requires pinhole placement over a finite aperture. Ding *et al*. [8] used a general linear camera (GLC) model for approximating forward projection for spherical mirrors. Agrawal *et al*. [1] proposed a forward projection equation for NCCS but only for the *axial* configuration, which restricts the positioning of the camera. They showed 3D reconstruction only using spherical mirrors, since any camera placement becomes an axial configuration for a spherical mirror. We present the analytical solution for off-axis case and show that the equation derived in [1] is a special case. Vandeportaele [29] also analyzed forward projection using a different approach involving intersection/tangency of quadrics. However, his approach results in a higher degree equation for several mirrors, both for off-axis camera placement and axial configuration, and also results in more complicated coefficients for the FP equation compared to ours.

Sparse 3D Reconstruction: Catadioptric cameras have been used for 3D reconstruction in previous works [14, 12, 17, 1, 20, 9, 22, 21]. In general, previous techniques have used a central approximation and minimized 3D or angular errors [19] instead of the image reprojection error. Micusik and Pajdla [20] and Lhuillier [18] used a central approximation for motion estimation followed by bundle-adjustment to minimize the reprojection error for sparse reconstruction. However, similar to [19], [20] used an optimization based approach for computing the mirror intersection point. Moreover, these approaches have only considered non-central axial configurations. Thus, previous approaches have not analyzed off-axis NCCS for 3D reconstruction, which we demonstrate. We avoid any central or GLC approximation and minimize the image reprojection error, thanks to our analytical projection model.

Dense Volumetric Reconstruction: Modification of stereo matching is often used for omnidirectional and panoramic images by reprojecting the images onto a different coordinate system [4, 13]. Ding *et al*. [8] proposed an epsilon-stereo constraint to allow slight vertical parallax between two non-central cameras (using spherical mirrors). However, as shown in [1], the epipolar curves for spherical mirror are quartic and thus the epsilon-stereo constraint does not model the vertical parallax correctly. In addition, a GLC approximation is used for fast forward projection, by tessellating the image into triangles and associating a GLC to each of them. Even when using the GLC approximation, forward projection requires a search in the image space: it needs to be checked if the projection of a 3D point lies within the triangle associated with that GLC projection. In contrast, our analytical projection model provides a one-shot solution to the projection of 3D points without any approximation.

Taguchi *et al*. [27] proposed an axial-cone model to enable fast back-projection of rays using GPU for NCCS in axial configuration. Using the axial-cone model, they apply plane sweeping [7, 16] for dense 3D reconstruction. However, when using multiple mirrors with a single camera, only a spherical mirror provides axial configuration for all the mirrors, which is a severe restriction. Our FP equation allows us to efficiently perform plane sweeping for multiple non-spherical mirrors for dense volumetric reconstruction. Kuthirummal and Nayar [15] proposed non-central radial imaging systems for 3D reconstruction. However, their setup is also axial. To the best of our knowledge, ours is the first paper to show sparse and dense 3D reconstruction using off-axis NCCS with multiple mirrors.

2. Forward Projection for NCCS

Consider a rotationally symmetric mirror described by the following equation in the world coordinate system:

$$x^2 + y^2 + Az^2 + Bz - C = 0. \tag{1}$$

Let the mirror axis be aligned with the z axis. On any plane containing the mirror axis, the mirror profile is given by a 2D conic $r^2 + Az^2 + Bz - C = 0$. This parametrization [25] describes spherical mirror $(A = 1, C + B^2/4 > 0)$ along with parabolic $(A = 0, C = 0)$, hyperbolic $(A < 0, C < 0)$

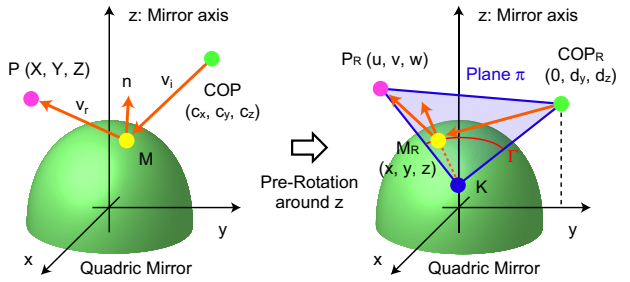

Figure 1. Use of pre-rotation and reflection plane π to simplify the derivation of FP equation.

and ellipsoidal ($A > 0, C > 0$) mirrors. Consider a pinhole camera placed at location $\texttt{COP} = [c_x, c_y, c_z]^T$ and a given 3D point $\mathbf{P} = [X, Y, Z]^T$ as shown in Figure 1 (left). The problem of forward projection is to find the point $\mathbf{M}$ on the mirror, such that ray joining COP with $\mathbf{M}$ after reflection passes through the given point $\mathbf{P}$. Note that unlike the derivation in [1], the analysis cannot be performed in a 2D plane containing the mirror axis.

Let $\mathbf{v}_i = \mathbf{M} - \texttt{COP}$ be the incoming ray, $\mathbf{n}$ be the normal of the mirror at $\mathbf{M}$, and $\mathbf{v}_r = \mathbf{P} - \mathbf{M}$ be the reflected ray. From the law of reflection, the following constraints are well-known: (a) Planarity: $\mathbf{v}_i$, $\mathbf{v}_r$, and $\mathbf{n}$ lie on the same plane, and (b) Angle Constraint: Angle between $\mathbf{v}_i$ and $\mathbf{n}$ is equal to the angle between $\mathbf{v}_r$ and $\mathbf{n}$. These two constraints result in the following reflection equation:

$$\mathbf{v}_r = \mathbf{v}_i - 2\mathbf{n}(\mathbf{v}_i^T\mathbf{n})/(\mathbf{n}^T\mathbf{n}). \tag{2}$$

In addition, since the point $\mathbf{M}$ lies on the mirror, it has to satisfy the mirror equation (1). We now present two key ideas to simplify the derivation of FP equation.

2.1. Pre-Rotation

Since the mirror is rotationally symmetric around the z axis, the entire coordinate system can be rotated around it to place the COP on the y axis, as shown in Figure 1 (right). Let $t = \sqrt{c_x^2 + c_y^2}$. Specifically, we compute the rotation matrix $R = \begin{bmatrix} c_y/t & -c_x/t & 0 \\ c_x/t & c_y/t & 0 \\ 0 & 0 & 1 \end{bmatrix}$ such that $R * \texttt{COP} = [0, d_y, d_z]^T$. By setting the x component of COP to be zero, we reduce the degree of subsequent equations, making it easier to solve. Note that this rotation is independent of the 3D point $\mathbf{P}$. Let $\mathbf{P}_R = R * \mathbf{P} = [u, v, w]^T$ be the new location of $\mathbf{P}$ after rotation. We compute the mirror intersection point $\mathbf{M}_R$ in this coordinate system and rotate it by R^{-1} to obtain the mirror intersection point in the original coordinate system.

2.2. Use of Reflection Plane π

Let $\mathbf{M}_R = [x, y, z]^T$. Let π denote the reflection plane on which $\mathbf{v}_i$, $\mathbf{v}_r$, and $\mathbf{n}$ lie. The normal at $\mathbf{M}_R$ can be computed from the mirror equation to be

$$\mathbf{n} = [x, y, Az + B/2]^T. \tag{3}$$

Since the mirror is rotationally symmetric around the z axis, the normal intersects the z axis at point $\mathbf{K} = [0, 0, z - Az - B/2]^T$, which also lies on π. Thus, the equation of π can be obtained using the points $\mathbf{K}$, $\texttt{COP}_R$ and $\mathbf{P}_R$ and is given by

$$c_1(z)x + c_2(z)y + c_3(z) = 0, \tag{4}$$

where

$$c_1(z) = (B + 2Az)(d_y - v) + 2d_y(w - z) + 2v(z - d_z),$$
$$c_2(z) = u(B + 2d_z - 2z + 2Az),$$
$$c_3(z) = ud_y(B + 2Az).$$

The key point to note is that the equation for π is *linear* in x and y. Using this equation, we can easily compute x in terms of y and z as

$$x = \frac{-c_2(z)y - c_3(z)}{c_1(z)}. \tag{5}$$

Substituting x in the mirror equation gives us our first intermediate equation $\mathbf{IE}_1$

$$\begin{aligned}\mathbf{IE}_1 : (c_1^2(z) + c_2^2(z))y^2 + 2c_2(z)c_3(z)y \\ + c_3^2(z) + c_1^2(z)(Az^2 + Bz - C) = 0.\end{aligned} \tag{6}$$

Note that this equation is *quadratic* in y, although the coefficients are functions of z. This equation describes the curve Γ given by the intersection of plane π with the mirror, on which reflection can happen. However, the correct point on this curve has to satisfy the angle constraint. Next, we show that by using the law of reflection, we obtain another intermediate equation $\mathbf{IE}_2$, which is also quadratic in y. By eliminating y between $\mathbf{IE}_1$ and $\mathbf{IE}_2$, we finally obtain a single 8^{th} degree equation in z.

2.3. Obtaining $\mathbf{IE}_2$ using Law of Reflection

To obtain $\mathbf{IE}_2$, we use the constraint that the reflected ray $\mathbf{v}_r$ should pass through the given point $\mathbf{P}_R$, yielding,

$$\mathbf{v}_r \times (\mathbf{P}_R - \mathbf{M}_R) = 0, \tag{7}$$

where $\times$ denotes the cross product. The incoming ray $\mathbf{v}_i$ is given by

$$\mathbf{v}_i = \mathbf{M}_R - \texttt{COP}_R = [x, y - d_y, z - d_z]^T. \tag{8}$$

By substituting $\mathbf{v}_i$ and $\mathbf{n}$ in the reflection equation (2), we obtain $\mathbf{v}_r$. Then $\mathbf{v}_r$ is substituted in (7) along with $\mathbf{P}_R$ and $\mathbf{M}_R$, resulting in the following three equations (since the cross product is in 3D):

$$\begin{aligned}E_1 &: k_{11}(z)x + k_{12}(z)y + k_{13}(z)xy + k_{14}(z)y^2 \\ &\quad + k_{15}(z) = 0. \\ E_2 &: k_{21}(z)x + k_{22}(z)y + k_{23}(z)xy + k_{24}(z) = 0. \\ E_3 &: k_{31}(z)y^2 + k_{32}(z)y + k_{33}(z) = 0.\end{aligned} \tag{9}$$

The coefficients $k_{ij}(z)$ are given in the supplementary materials. We note that the three equations in (9) are not independent. Substituting x using (5) in any of these three equations leads to the *same* equation. Thus, these are dependent equations and any one of them can be chosen for $\mathbf{IE}_2$. Notice that E_3 is independent of x, and thus we choose it to be $\mathbf{IE}_2$. If pre-rotation is not performed, E_1, E_2, and E_3 will have higher degree terms in x and y, making it more difficult to solve and will lead to a higher degree FP equation.

One can also show that the equation of plane π can be obtained by manipulating these three equations as follows: (a) Eliminate y^2 from E_1 and E_3 to get equation $E_4 \equiv k_{14}(z)E_3 - k_{31}(z)E_1$, and (b) Eliminate xy from E_2 and E_4. However, this is non-obvious. Thus, use of plane π simplifies the derivation.

2.4. Obtaining Forward Projection Equation

Now we eliminate y from $\mathbf{IE}_1$ and $\mathbf{IE}_2$ to obtain a single equation in only one unknown z. Notice that there are no x terms in both $\mathbf{IE}_1$ and $\mathbf{IE}_2$. We first rewrite $\mathbf{IE}_1$ and $\mathbf{IE}_2$ as

$$\begin{aligned} \mathbf{IE}_1 &: k_{41}(z)y^2 + k_{42}(z)y + k_{43}(z) = 0, \\ \mathbf{IE}_2 &: k_{31}(z)y^2 + k_{32}(z)y + k_{33}(z) = 0, \end{aligned} \tag{10}$$

where $k_{41}(z) = c_1^2(z) + c_2^2(z)$, $k_{42}(z) = 2c_2(z)c_3(z)$ and $k_{43}(z) = c_3^2(z) + c_1^2(z)(Az^2 + Bz - C)$. Eliminating y^2, we get

$$y = -\frac{k_{41}(z)k_{33}(z) - k_{31}(z)k_{43}(z)}{k_{41}(z)k_{32}(z) - k_{31}(z)k_{42}}. \tag{11}$$

Substituting y back into either $\mathbf{IE}_1$ or $\mathbf{IE}_2$ gives the desired forward projection equation:

$$\begin{aligned} &k_{41}(z)\left(k_{43}(z)k_{32}^2(z) - k_{42}(z)k_{32}(z)k_{33}(z) + k_{41}(z)k_{33}^2(z)\right) \\ &- k_{31}(z)(-k_{33}(z)k_{42}^2(z) + k_{43}(z)k_{32}(z)k_{42}(z) \\ &+ 2k_{41}(z)k_{43}(z)k_{33}(z)) + k_{43}^2(z)k_{31}^2(z) = 0. \end{aligned}$$

After simplification[1], this turns out to be an 8^{th} degree polynomial equation in z. The coefficients depend on the known mirror parameters (A, B, C), the known location of COP (d_y, d_z) and the known 3D point (u, v, w).

By setting $d_y = 0$, the NCCS reduces to the axial configuration. In that case, we obtain a 6^{th} degree equation equivalent to the equation derived in [1]. Thus, the FP equation derived in [1] for axial NCCS is a special case of our FP equation for general off-axis NCCS.

3. Analysis of Forward Projection Equation

Table 1 lists the degree of FP equation for various mirror shapes and camera placement. SVP refers to single-viewpoint configuration, while NSVP refers to non-single viewpoint configuration. A $-$ indicates that SVP constraint may be impossible or unpractical to achieve with that particular mirror.

[1]Supplementary materials contain Matlab code deriving this equation along with intermediate steps.

Table 1. Degree of forward projection equation for various quadric mirrors and camera placement. A $-$ indicates that single-viewpoint constraint (SVP) may be impossible or unpractical to achieve. NSVP refers to non-single viewpoint.

Mirror Shape	Parameters	Camera Placement		
		Off-Axis	Axial	
			NSVP	SVP
General	A, B, C	8	6	-
Spherical	$A = 1, B = 0, C > 0$		$4^{\ddagger}$	-
Elliptical	$A > 0, B = 0, C > 0$	8	6	2
Hyperbolic	$A < 0, B = 0, C < 0$	8	6	2
Parabolic	$A = 0, C = 0$	7	5	$2^{\dagger}$
Conical	$A < 0, B = 0, C = 0$	4	2	-
Cylindrical	$A = 0, B = 0, C > 0$	4	2	-

‡ For a spherical mirror, any camera placement is axial configuration with NSVP.
† Orthographic camera placed on the mirror axis.

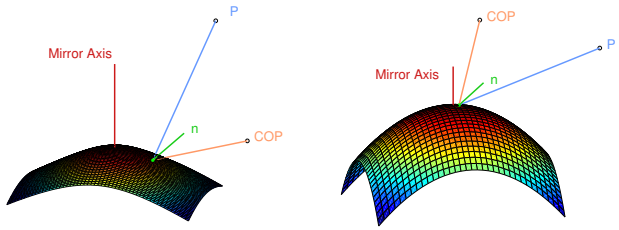

Figure 2. Visualization of forward projection for a hyperbolic mirror (left) and a parabolic mirror (right).

For a general quadric and off-axis camera placement, the FP equation has a degree of 8. For a parabolic mirror $(A = 0, C = 0)$, the degree reduces to 7. For a cylindrical mirror polished on outside $(A = 0, B = 0, C > 0)$, the degree reduces to 4 for off-axis camera placement. Axial-configuration with cylindrical mirror is practical when the mirror is polished on the inside [15]. The degree of FP equation in that case reduces to 2. Figure 2 shows an example of computed mirror intersection point along with incoming and reflected rays for hyperbolic and parabolic mirrors.

Note that when $A \neq 0$, B can be set to 0 by shifting the mirror along the z axis. Table 1 also shows the degree of FP equation for conical $(C = 0, A < 0)$, hyperbolic $(A < 0, C < 0)$, and ellipsoidal $(A > 0, C > 0)$ mirrors. For mirrors and camera placement satisfying the single-viewpoint constraint (SVP), the degree is 2. This is obvious, since the mirror intersection point can be obtained by finding the intersection of the mirror with the ray joining the given 3D point and the effective (virtual) center of projection. For other radial imaging systems [15] using mirrors polished on inside $(A < 0, C > 0)$, the degree of forward projection equation is 8 and 6 for off-axial and axial camera placement respectively. We note that [15] describes epipolar geometry and 3D reconstruction for radial imaging systems, but does not derive the forward projection equation.

In contrast, the approach in [29] using intersection and

tangency of quadrics results in higher degree equations for both off-axis camera placement and axial configuration. For off-axis camera placement, [29] reports 6^{th} degree equation for cylindrical mirror and 8^{th} degree equation for conical mirror, compared to 4^{th} degree equation we obtain for both of these mirrors. For axial configuration, [29] reports 6^{th} degree equation for conical mirror and 8^{th} degree equation for ellipsoidal mirror, compared to 2^{nd} degree equation and 6^{th} degree equation we obtain respectively.

Numerical Accuracy and Run-Time: To solve the FP equation, we use the Matlab `roots` command, which computes the eigenvalues of the corresponding companion matrix. The correct solution is found by checking the Snell's law for each real solution. Using a hyperbolic mirror with parameters $A = -1.2$, $B = 3.4$ and $C = -33.2$, we performed 10^5 trials by choosing random locations for the COP and 3D point P. The median error in angle constraint was of the order of 10^{-10}. Thus, our FP equation is numerically stable to compute the mirror intersection point. Projection of 10^5 points took 41.4 seconds using our analytical FP equation compared to 1695.6 seconds using IFP in Matlab on a standard PC. Thus, our analytical FP equation offers a 40 times speed up over iterative forward projection. For 3D reconstruction, this amounts to reducing the run-time from approximately two days to one hour.

4. 3D Reconstruction using NCCS

Sparse and dense 3D reconstruction using perspective cameras is a well-studied problem in computer vision and has matured significantly over the last two decades. However, 3D reconstruction using non-central catadioptric cameras is considered challenging, due to the lack of exact models. Previous work has been mostly restricted to axial configurations, *e.g.*, a perspective camera placed on the axis of a spherical or hyperbolic mirror [20]. We now show that using our FP equation, sparse and dense 3D reconstruction can be done efficiently for the general case of off-axis camera placement.

4.1. Sparse 3D Reconstruction in Simulations

We first describe sparse 3D reconstruction using bundle adjustment (BA) in simulations. We choose a setup of a single perspective camera looking at four identical parabolic mirrors. The equation of each mirror is given by $x^2+y^2+40z = 0$. The vertices of the four parabolic mirrors are placed at $[\pm 20, \pm 20, 200]$ mm in the camera coordinate system and their axes are aligned to the z axis; thus the camera is significantly off-axis for all the mirrors. The internal camera calibration and mirror parameters are assumed to be known. We randomly place 3D points in a hemisphere centered at $[0, 0, 200]$ with a radius of 1000 mm, and compute the ground truth projections for all points and for all mirrors using our FP equation. Gaussian noise ($\sigma = [0 - 1]$ pixels) is then added to the ground truth image projections, as well as location ($\sigma = 0.5$ mm) and orientation ($\sigma = 1^\circ$) of all the mirrors. Initial reprojection error was 5% of the image size and the initial 3D RMSE error was ~ 450 mm. The bundle adjustment optimization involves estimating the calibration (location and axis of mirror) along with the location of 3D points from noisy observations and the initial calibration estimates. The initial estimates of the 3D points are obtained by back-projecting corresponding rays and finding their intersection. Then the image reprojection error is minimized.

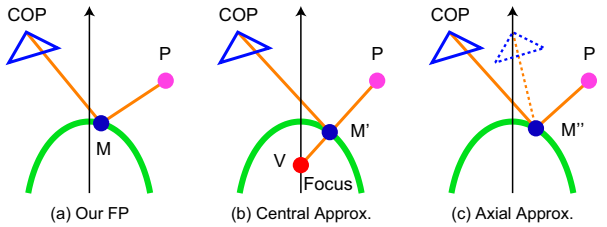

Figure 3. (a) Our FP equation computes the exact mirror reflection point M given the COP and the 3D point P. (b) Achieving central approximation by using the focus of the parabola as the virtual projection center V and computing M'. (c) Achieving axial approximation by placing the COP on the mirror axis to compute M'' using [1].

We compare the reconstruction errors using (a) our FP equation, (b) central approximation, (c) axial approximation, and (d) iterative forward projection (IFP). Figure 3 shows how we compute the mirror reflection point using central and axial approximations. For IFP, we first use the central approximation to compute the initial estimate of the image projection and then refine it by minimizing the distance between the 3D point and the back-projected ray using non-linear optimization. Figure 4 shows estimation errors for different image noise levels. Note that the small mirror pose perturbation (location $\sigma = 0.5$ mm and orientation $\sigma = 1^\circ$ in this experiment) severely affects the 3D reconstruction performance (shown as initial errors). Both central and axial approximations fail to correct the errors in calibration and 3D point locations. In contrast, our analytical FP and IFP show the same estimation performance; however, the average run time of bundle adjustment using our FP was only 13 seconds, while that of IFP was 560 seconds, since IFP requires iterative optimization for each forward projection. Thus, we achieve a speed up of ~ 40X.

4.2. Sparse 3D Reconstruction using Images

We now analyze sparse 3D reconstruction results using synthetic images generated in POV-Ray as well as real dataset (Figure 5). Our reconstruction procedure consists of the following three steps: (a) estimate initial pose of all mirrors, (b) compute corresponding feature points across mirror images, and (c) iterate bundle adjustment with outlier removal. The initial mirror pose is estimated by marking sev-

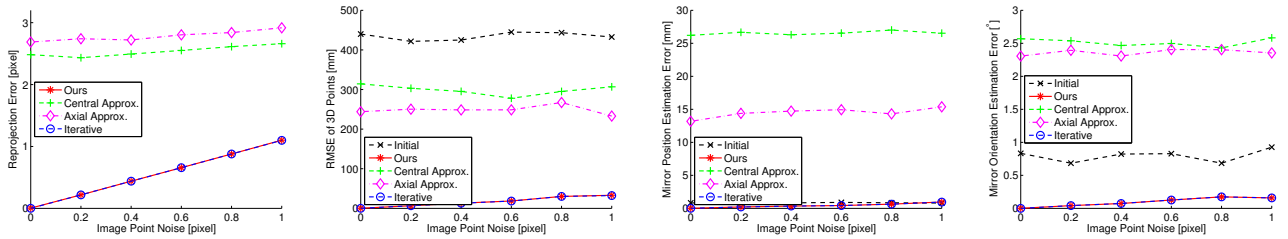

Figure 4. Bundle adjustment simulations using four parabolic mirrors and 100 3D points for different image noise levels. From left to right: Reprojection error, RMSE of reconstructed 3D points, mirror position estimation error, and mirror orientation estimation error. Both central and axial approximation fail to correctly minimize the reprojection error as well as error in reconstructed 3D points.

Figure 5. (Left) Synthetic image rendered with POV-Ray and zoom of a mirror image. (Right) Real captured photo and zoom of a mirror image. SIFT feature correspondences across mirror images are superimposed on the zoomed in mirror images. Red dots and green crosses respectively represent inliers and outliers determined by our iterative bundle adjustment procedure.

eral points on the mirror boundary and computing the corresponding rays in 3D. Since the mirror section boundary is a 2D circle with a known radius, its position/orientation can be computed in 3D such that these rays pass through the circle. Correspondences across mirror images are obtained using SIFT features. Finally, we iterate bundle adjustment by removing outliers of the SIFT correspondences if the reprojection error of the reconstructed 3D point is greater than twice the average reprojection error as in [1].

Result using POV-Ray: Figure 5 (left) shows a 2000×2000 pixels rendered image of four parabolic mirrors placed at $[\pm 15, \pm 15, 200]$ mm in the camera coordinate system. The mirrors are surrounded by a cube of 1000 mm, whose walls are textured. Instead of using the ground truth values for each mirror pose, we estimated it as described above. This introduced a position estimation error of 0.14 mm and an orientation estimation error of $0.94°$ (averaged over four mirrors). In Section 4.3, we demonstrate that this calibration error can result in significant errors in subsequent dense volumetric reconstruction. Figure 6 demonstrates that our iterative bundle adjustment process reconstructs the correct 3D points by refining the mirror pose estimates. By using the refined calibration after bundle adjustment, the dense volumetric reconstruction is improved as shown in Section 4.3.

Result using Real Setup: We captured a 22 megapixel photo of four parabolic mirrors ($x^2 + y^2 + 40z = 0$, cross-section radius 20 mm) with an interval of $\sim$ 80 mm (Figure 5 (right)) using a Mamiya 645AFD camera. The resolution of each mirror image was $\sim$ 2M pixels. The mirrors are cut using a milling machine and have surface artifacts, leading to jagged edges in the image. The distance between the camera and the mirrors was $\sim$ 480 mm. The initial mirror poses were computed using the boundary-based approach described above. Figure 7 (right) shows the reconstructed 3D points along with their corresponding fitted planes. Again note that before bundle adjustment, the 3D points are incorrectly reconstructed near the mirrors.

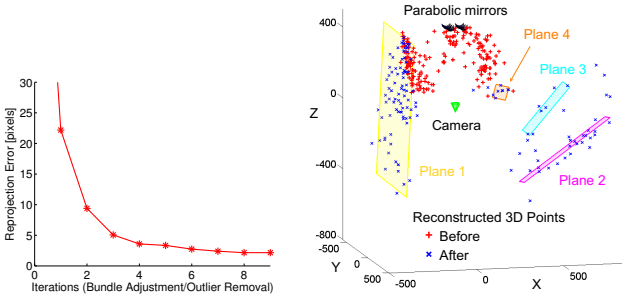

Figure 7. Sparse 3D reconstruction results for the real image. (Left) Reprojection error with BA iterations. Initial reprojection error was 87 pixels for 22 megapixel image. (Right) Visualization of the reconstructed 3D points before and after BA.

4.3. Dense Volumetric 3D Reconstruction

Ding *et al*. [8] used a GLC approximation to perform fast forward projection for dense reconstruction with spherical

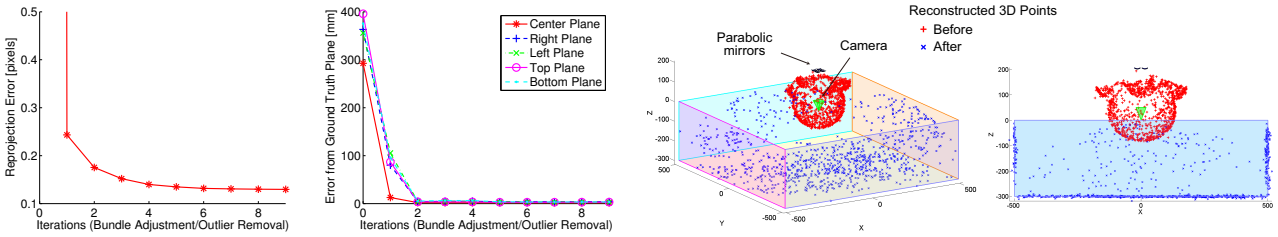

Figure 6. Sparse 3D reconstruction results for synthetic image generated using POV-Ray. (Left) Reprojection error decreases with the bundle adjustment (BA) and outlier removal iterations. Initial reprojection error was 32 pixels for 4 megapixel image. (Middle-Left) Average distance error of reconstructed 3D points from the ground truth planes shown in Figure 5. (Right) Visualization of the reconstructed 3D points before and after BA from two different viewpoints. Before BA, the 3D points are incorrectly reconstructed near the mirrors.

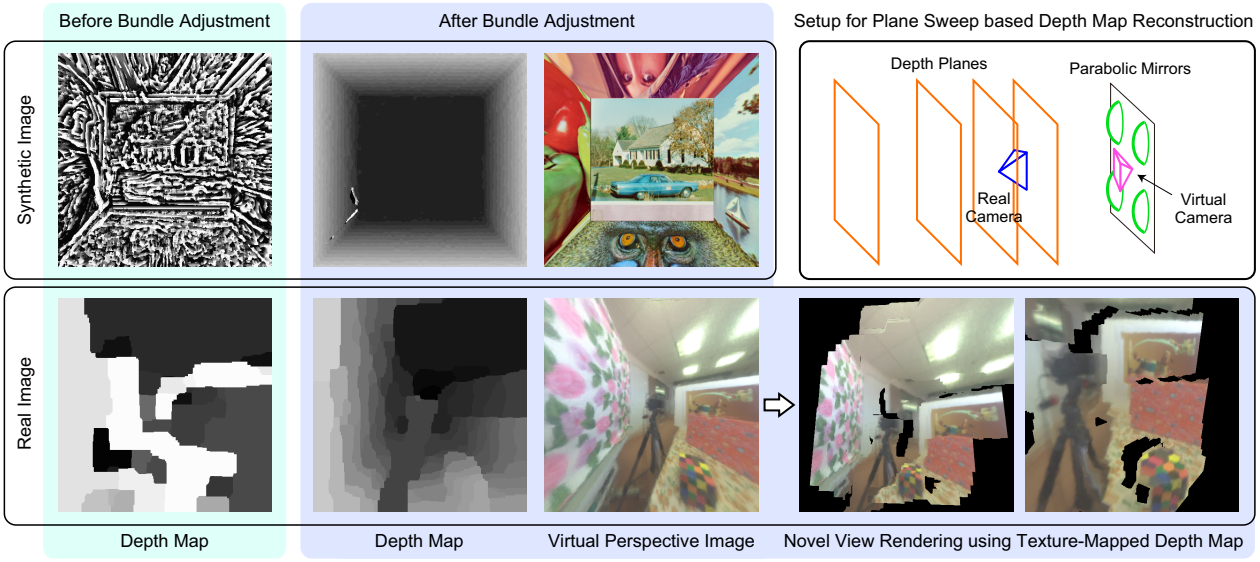

Figure 8. Using the refined calibration after BA significantly improves the dense depth map compared to that obtained using the initial calibration before BA. For synthetic image, the estimated depth map is close to ground truth after BA, even without any regularization/smoothing. FOV of the depth maps and virtual perspective images is $120^\circ \times 120^\circ$. Bottom-right show the depth map rendered from two novel viewpoints and textured using the virtual perspective image. Black pixels correspond to regions that are occluded or are outside the FOV. All the images in this figure are horizontally flipped to visually match the scene layout with the input images.

mirror array. Our analytical FP equation avoids GLC approximation and enables us to perform fast volumetric 3D reconstruction using plane sweeping [7, 16]. As shown in Figure 8, we place a virtual camera at the center of the four parabolic mirrors and define a set of depth planes in the coordinate system of the virtual camera. For each mirror, we compute a set of layer images by computing the image projection of all the 3D points on the depth planes. This produces four images corresponding to four mirrors for each depth plane. We compute normalized cross correlation between them as the matching cost. Averaging the four images for each depth plane produces refocusing effect [27].

Figure 8 shows depth maps computed using the mirror pose estimates before and after bundle adjustment. We generated the depth maps using the resolution of 400×400 pixels, and using 20 and 40 depth planes for the synthetic and real images, respectively. The depth map of the synthetic image was generated by simply selecting the depth having the best matching cost *without* smoothing. For the real image, we regularized the matching cost using standard graph cuts approach used for stereo matching. Similar to sparse 3D reconstruction case, the initial mirror poses produce large errors in the estimated depth maps. Bundle adjustment refines the mirror pose estimates and enables to produce high-quality depth maps. A virtual perspective image from the virtual viewpoint is generated using the depth map; this corresponds to distortion correction of non-central images using scene geometry [24]. Figure 8 also shows novel viewpoint rendering using the virtual perspective image and texture mapped depth map from two novel viewpoints.

5. Conclusions

By deriving an analytical projection model for non-central catadioptric cameras, our paper bridges the gap between perspective and catadioptric cameras for applications

such as 3D reconstruction. It makes non-central catadioptric cameras as attractive and easy to use as perspective cameras. We demonstrated a complete pipeline for 3D reconstruction using NCCS: feature matching and bundle adjustment for sparse 3D reconstruction by minimizing the image reprojection error, followed by dense volumetric reconstruction. Several algorithms developed for perspective cameras can be directly applied to NCCS by replacing the perspective projection equation with the derived FP equation. We hope that the our paper will stimulate further research in modeling and analysis of catadioptric imaging systems along with designing better image features tailored for catadioptric images. Computing exact light path via quadric mirror reflection may be useful in other applications such as rendering of specular surfaces, specular shape estimation, as well as modeling and simulation of general mirror based imaging systems.

References

[1] A. Agrawal, Y. Taguchi, and S. Ramalingam. Analytical forward projection for axial non-central dioptric and catadioptric cameras. In *ECCV*, volume 6313, pages 129–143, 2010.

[2] M. Baker. Alhazen's problem. *American J. Mathematics*, 4(1):327–331, 1881.

[3] S. Baker and S. Nayar. A theory of single-viewpoint catadioptric image formation. *IJCV*, 35(2):175–196, Nov. 1999.

[4] R. Bunschoten and B. Kröse. Robust scene reconstruction from an omnidirectional vision system. *IEEE Trans. Robotics and Automation*, 19(2):351–357, Apr. 2003.

[5] V. Caglioti, P. Taddei, G. Boracchi, S. Gasparini, and A. Giusti. Single-image calibration of off-axis catadioptric cameras using lines. In *OMNIVIS*, Oct. 2007.

[6] M. Chen and J. Arvo. Theory and application of specular path perturbation. *ACM Trans. Graph.*, 19(4):246–278, 2000.

[7] R. Collins. A space-sweep approach to true multi-image matching. In *CVPR*, pages 358–363, Jun 1996.

[8] Y. Ding, J. Yu, and P. Sturm. Multi-perspective stereo matching and volumetric reconstruction. In *ICCV*, pages 1827–1834, Sept. 2009.

[9] C. Geyer and K. Daniilidis. Structure and motion from uncalibrated catadioptric views. In *CVPR*, pages 279–286, 2001.

[10] G. Glaeser. Reflections on spheres and cylinders of revolution. *J. Geometry and Graphics*, 3(2):121–139, 1999.

[11] N. Gonçalvez and A. C. Nogueira. Projection through quadric mirrors made faster. In *OMNIVIS*, pages 2141–2148, 2009.

[12] M. Kanbara, N. Ukita, M. Kidode, and N. Yokoya. 3D scene reconstruction from reflection images in a spherical mirror. In *ICPR*, volume 4, pages 874–879, 2006.

[13] S. B. Kang and R. Szeliski. 3-D scene data recovery using omnidirectional multibaseline stereo. In *CVPR*, pages 364–370, 1996.

[14] Y. Kojima, R. Sagawa, T. Echigo, and Y. Yagi. Calibration and performance evaluation of omnidirectional sensor with compound spherical mirrors. In *OMNIVIS*, 2005.

[15] S. Kuthirummal and S. Nayar. Multiview radial catadioptric imaging for scene capture. *ACM Trans. Graph.*, 25(3):916–923, 2006.

[16] K. Kutulakos and S. Seitz. A theory of shape by space carving. *IJCV*, 38(3):199–218, 2000.

[17] D. Lanman, D. Crispell, M. Wachs, and G. Taubin. Spherical catadioptric arrays: Construction, multi-view geometry, and calibration. In *3DPVT*, pages 81–88, 2006.

[18] M. Lhuillier. Toward flexible 3d modeling using a catadioptric camera. In *CVPR*, pages 1–8, June 2007.

[19] M. Lhuillier. Automatic scene structure and camera motion using a catadioptric system. *Computer Vision and Image Understanding*, 109(2):186–203, Feb. 2008.

[20] B. Micusik and T. Pajdla. Autocalibration and 3D reconstruction with non-central catadioptric cameras. In *CVPR*, pages 58–65, 2004.

[21] S. Nayar. Sphereo: Determining depth using two specular spheres and a single camera. In *SPIE Conf. Optics, Illumination, and Image Sensing for Machine Vision III*, pages 245–254, Nov. 1988.

[22] S. A. Nene and S. K. Nayar. Stereo with mirrors. In *ICCV*, pages 1087–1094, Jan. 1998.

[23] S. Roth and M. Black. Specular flow and the recovery of surface structure. In *CVPR*, volume 2, pages 1869–1876, June 2006.

[24] R. Swaminathan, M. Grossberg, and S. Nayar. A perspective on distortions. In *CVPR*, volume 2, pages 594–601, 2003.

[25] R. Swaminathan, M. Grossberg, and S. Nayar. Non-single viewpoint catadioptric cameras: Geometry and analysis. *IJCV*, 66(3):211–229, Mar 2006.

[26] Y. Taguchi, A. Agrawal, S. Ramalingam, and A. Veeraraghavan. Axial light field for curved mirrors: Reflect your perspective, widen your view. In *CVPR*, pages 499–506, June 2010.

[27] Y. Taguchi, A. Agrawal, A. Veeraraghavan, S. Ramalingam, and R. Raskar. Axial-cones: Modeling spherical catadioptric cameras for wide-angle light field rendering. *ACM Trans. Graph.*, 29(6):172:1–172:8, Dec. 2010.

[28] J. Unger, A. Wenger, T. Hawkins, A. Gardner, and P. Debevec. Capturing and rendering with incident light fields. In *EGSR*, pages 141–149, 2003.

[29] B. Vandeportaele. *Contributions à la vision omnidirectionnelle: Étude, Conception et Étalonnage de capteurs pour lacquisition dimages et la modélisation 3D*. PhD thesis, Institut National Polytechnique de Toulouse, France, Dec. 2006. in French.