



May 2002

Paracatadioptric Camera Calibration

Christopher Geyer
University of Pennsylvania

Kostas Daniilidis
University of Pennsylvania, kostas@cis.upenn.edu

Follow this and additional works at: http://repository.upenn.edu/cis_papers

Recommended Citation

Christopher Geyer and Kostas Daniilidis, "Paracatadioptric Camera Calibration", . May 2002.

Copyright 2002 IEEE. Reprinted from *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Volume 24, Issue 5, May 2002, pages 687-695.
Publisher URL: <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isNumber=21601&puNumber=34>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/cis_papers/56
For more information, please contact libraryrepository@pobox.upenn.edu.

Paracatadioptric Camera Calibration

Abstract

Catadioptric sensors refer to the combination of lens-based devices and reflective surfaces. These systems are useful because they may have a field of view which is greater than hemispherical, providing the ability to simultaneously view in any direction. Configurations which have a unique effective viewpoint are of primary interest, among these is the case where the reflective surface is a parabolic mirror and the camera is such that it induces an orthographic projection and which we call paracatadioptric. We present an algorithm for the calibration of such a device using only the images of lines in space. In fact, we show that we may obtain all of the intrinsic parameters from the images of only *three* lines and that this is possible without any metric information. We propose a closed-form solution for focal length, image center, and aspect ratio for skewless cameras and a polynomial root solution in the presence of skew. We also give a method for determining the orientation of a plane containing two sets of parallel lines from one uncalibrated view. Such an orientation recovery enables a rectification which is impossible to achieve in the case of a single uncalibrated view taken by a conventional camera. We study the performance of the algorithm in simulated setups and compare results on real images with an approach based on the image of the mirror's bounding circle.

Keywords

Omnidirectional vision, panoramic vision, catadioptric camera, vanishing points, calibration

Comments

Copyright 2002 IEEE. Reprinted from *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Volume 24, Issue 5, May 2002, pages 687-695.

Publisher URL: <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isNumber=21601&puNumber=34>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Short Papers

Paracatadioptric Camera Calibration

Christopher Geyer, *Member, IEEE*, and
Kostas Daniilidis, *Member, IEEE*

Abstract—Catadioptric sensors refer to the combination of lens-based devices and reflective surfaces. These systems are useful because they may have a field of view which is greater than hemispherical, providing the ability to simultaneously view in any direction. Configurations which have a unique effective viewpoint are of primary interest, among these is the case where the reflective surface is a parabolic mirror and the camera is such that it induces an orthographic projection and which we call paracatadioptric. We present an algorithm for the calibration of such a device using only the images of lines in space. In fact, we show that we may obtain all of the intrinsic parameters from the images of only *three* lines and that this is possible without any metric information. We propose a closed-form solution for focal length, image center, and aspect ratio for skewless cameras and a polynomial root solution in the presence of skew. We also give a method for determining the orientation of a plane containing two sets of parallel lines from one uncalibrated view. Such an orientation recovery enables a rectification which is impossible to achieve in the case of a single uncalibrated view taken by a conventional camera. We study the performance of the algorithm in simulated setups and compare results on real images with an approach based on the image of the mirror's bounding circle.

Index Terms—Omnidirectional vision, panoramic vision, catadioptric camera, vanishing points, calibration.

1 INTRODUCTION

A catadioptric instrument is an optical system combining reflective (catoptric) and refractive (dioptric) elements [1]. Catadioptric combinations have been extensively used in telescopes in order to focus light from the stars onto the eye of the observer. Recently, they have been introduced in robotics and computer vision to enable omnidirectional sensing.

We classify the catadioptric systems in two groups, central and noncentral, based on the uniqueness of an effective viewpoint. The focal properties of mirrors with a conic profile were discovered by Diocles [2]. Nayar [3], Baker and Nayar [4] gave the first formal treatment of catadioptric systems with a single viewpoint in the context of computer vision and their geometry has been also studied in [5], [6], [7]. Uniqueness of an effective viewpoint is desirable because it allows the mapping of any part of the scene to a perspective plane exactly as if it were taken with a perspective camera whose focus is the effective viewpoint. In this sense, a central catadioptric system has the same effect as a camera rotating about its focus. Furthermore, easily modified multiple view algorithms can be applied for reconstruction [8], [9]. Central catadioptric systems are extensively used now for visualization [10], [11] and navigation [12], [13], [14]. Pyramidal multifaceted mirrors mounted above clusters of cameras can simultaneously achieve high-resolution and one effective viewpoint [15]. For a broad coverage of central, as well as, noncentral omnidirectional sensors the reader is referred to an extensive review by Yagi [16] as well as to the proceedings of the Workshop for Omnidirectional vision [17] and to the collection in [18].

- The authors are with GRASP Laboratory, University of Pennsylvania, 3401 Walnut St., 336C, Philadelphia, PA 19104-6228. E-mail: {kostas, geyer}@grasp.cis.upenn.edu.

Manuscript received 22 Mar. 2000; revised 16 Nov. 2000; accepted 29 Aug. 2001.

Recommended for acceptance by R. Kumar.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 111747.

This paper deals with central catadioptric sensors of paraboloid reflective and orthographic refractive components. We call these sensors, **paracatadioptric**. Our contribution is an algorithm for calibrating a paracatadioptric camera using a single view of three lines. By calibration, we mean the estimation of the mirror center, the combined focal length of the lens and the mirror, and the aspect ratio and skew of the CCD-chip. No metric assumptions are made about the scene or the construction and the image of the mirror. Aspect ratio, effective focal length, and mirror center can be found in closed form from the images of $N \geq 3$ lines in space if the camera skew is zero. In the presence of skew, the solutions for aspect ratio and skew are polynomial roots. After solving for them, focal length and mirror center can be found in closed form. We remark that perspective cameras can not be calibrated from a single frame without metric information.

Knowledge of the calibration parameters enables the mapping from pixels to rays through the unique effective viewpoint. This mapping can be used for unwarping a portion of the original image or for any stereo or motion algorithm.

We further show that by adding affine information—two sets of parallel lines—we can estimate the attitude of the plane spanned by the directions of the two parallel lines. This is also possible in calibrated perspective cameras, however, our solution provides a geometric illustration of the line at infinity and the vanishing points in paracatadioptric projection. Knowledge of the attitude of a plane enables the rectification of the plane from one view with only affine information.

Without knowledge of the projective properties of parabolic projection, one might consider two possible calibration algorithms. The most obvious would be to fit a circle to the image of the boundary of the mirror [19], [20]. Typically, these devices are designed so that the mirror's boundary is a circle visible within the image. Then, assuming that the mirror is constructed properly and that the field of view is known, it is a simple matter of calculating the focal length, for then the center of the circle is the mirror center. This approach has some shortcomings: it assumes the mirror boundary accurately encodes the intrinsics and that the field of view is known or previously calibrated, and it requires the visibility of the mirror boundary. However, it is advantageous in that it can be easily automated and does not require the use of line images which our algorithm does. We compare this possible algorithm with the one proposed in this article.

The outline of the paper is as follows: We describe the geometric properties of paracatadioptric cameras in Section 2. In Section 3, we describe the algorithms for calibration and rectification, and in Section 4 we present results in real images as well as in simulations.

2 PROJECTIVE PROPERTIES

In this section, we develop projective and invariant properties which we will later find useful in developing our calibration algorithm.

We assume that we have a mirror in the shape of a paraboloid with focal length f_0 (cm) whose axis of symmetry is parallel to the direction of an orthographic imaging device with scale s (pixels/cm) and aspect ratio of one. If $f = sf_0$, then, centimeters cancel and we have a measurement of combined focal length in pixels. For calibration, it is not necessary to estimate either s or f_0 and, in fact, it is impossible to determine one without knowledge of the other, even with metric information in the scene. Thus, we only estimate f . Since an orthographic projection is used, the distance from the paraboloid to the image plane and any translation perpendicular to the plane is irrelevant (the latter only inducing a translation in the image plane). We assume without loss of generality that the image plane is the plane of the directrix of the paraboloid. Recall that the directrix defines the plane: a point on the paraboloid is equidistant to the plane of the directrix and the paraboloid's focus.




Fig. 1. A cross-section of the paraboloid. The parabolic projection of a point P in space is obtained by orthographically projecting to the directrix the intersection of FP with the parabola.

Definition 1. The parabolic projection P' of a point P is the orthographic projection of the intersection Q' of the ray FP with the paraboloid to the plane of the directrix. Here, F is the focus of the paraboloid. (See Fig. 1.)

It can be shown, by intersecting a parameterized ray with the implicit equation of a paraboloid, then orthographically projecting to the image plane, that the image of a point (x, y, z) is given by

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \frac{2fx}{\sqrt{x^2+y^2+z^2}-z} \\ \frac{2fy}{\sqrt{x^2+y^2+z^2}-z} \end{pmatrix}, \quad (1)$$

except when $x = y = 0$ and $z \geq 0$ in which case the projection is undefined. When measured from a CCD chip, the point may be transformed by

$$\begin{pmatrix} u' \\ v' \end{pmatrix} = \begin{pmatrix} \alpha & \beta \\ 0 & \alpha^{-1} \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} + \xi. \quad (2)$$

It will be our aim to find $\xi = (\xi_x, \xi_y)^T$ (mirror center), f , $a = \alpha^2$ (aspect ratio), and β (skew).

In the following proposition, we determine the images of lines and their invariants. That images of lines are circles is mentioned by Nene and Nayar [21].

Proposition 1. A line in space is mapped to an arc of a circle, unless it intersects the optical axis, in which case it is mapped to a line; such a projection is called a **line image**. When the image is a part of a circle, the circle's radius r and the distance d from its center to the mirror center satisfy

$$4f^2 = r^2 - d^2. \quad (3)$$

Proof. The implicit equation of a parabola with focal length f is

$$\frac{1}{4f}(x^2 + y^2) - f = z. \quad (4)$$

Let P be the plane containing the focus F and the line in space. The orthographic projection of the intersection of the plane and the parabola contains the image of the line. Assume that the plane has a normal equal to $\hat{n} = (n_x, n_y, n_z)$, where $\|\hat{n}\| = 1$.

If $n_z = 0$, then the plane contains the optical axis and is also perpendicular to the image plane, thus the image of this line is a line in the image plane through $(0, 0)$ and (n_x, n_y) . In this case, (3) is not satisfied because r and d are undefined.

If, however, $n_z \neq 0$, then the plane has the implicit equation

$$n_x x + n_y y + n_z z = 0,$$

since F is the origin. Solving for z and substituting (4) we find that the orthographically projected points satisfy,

$$n_x x + n_y y + n_z \left(\frac{x^2 + y^2}{4f} - f \right) = 0.$$

We multiply this equation by $\frac{4f}{n_z}$, substitute $n_x^2 + n_y^2 = 1 - n_z^2$, factor and find that,

$$\left(x + 2f \frac{n_x}{n_z} \right)^2 + \left(y + 2f \frac{n_y}{n_z} \right)^2 = \left(\frac{2f}{n_z} \right)^2, \quad (5)$$

which is the equation of a circle.

Now to show that the invariant equation is satisfied,

$$d^2 = c_x^2 + c_y^2 = 4f^2 \frac{n_x^2}{n_z^2} + 4f^2 \frac{n_y^2}{n_z^2}.$$

Thus,

$$r^2 - d^2 = 4f^2 \frac{1}{n_z^2} - 4f^2 \frac{n_x^2}{n_z^2} - 4f^2 \frac{n_y^2}{n_z^2} = 4f^2, \quad \square$$

We remark that condition (4) is equivalent to the condition that a line image intersect the fronto-parallel horizon antipodally. The fronto-parallel horizon is the projection of the plane parallel to the image plane (plane of the directrix) and is a circle centered at the mirror center with radius $2f$.

Corollary 1. In a parabolic projection, where the mirror center is at the origin and the focal length is f , a sphere whose equator is a line image contains the points $(0, 0, -2f)$ and $(0, 0, 2f)$.

Proof. A sphere whose equator is a line image has an implicit equation

$$(c_x - x)^2 + (c_y - y)^2 + z^2 = r^2.$$

Substitute $(x, y, z) = (0, 0, \pm 2f)$, and one obtains the invariant (3) satisfied by any line image. \square

We use without proof the following two facts: First, any affine transformation (in particular, a change of aspect ratio or skew transformation) individually transforms a family of one or more circles to a family of ellipses with parallel axes and identical eccentricity. Second, there is a single affine transformation, unique up to scale, translation and rotation, which maps such a family of ellipses to a family of circles.

3 CALIBRATION AND RECTIFICATION

In this section, we present an algorithm for full calibration of a parabolic projection. This includes determining the mirror center ξ , focal length f , aspect ratio a (the ratio of independent scales in each axis) and skew β . We show that there are explicit formulas for all except when skew is unknown, in which case, all but skew and aspect ratio are explicit.

First, we recall [22], [23] a least-squares method for fitting a circle to a set of points. We minimize the sum of squares of the difference of square distances to the circle; we use this algebraic distance [24] as opposed to the sum of geometric distances because the geometric distance involves a square root which would require nonlinear minimization and has been shown in [22] to be more sensitive to outliers. The residual of a circle with radius r and center \mathbf{c} fitted to a set of points $\{\mathbf{p}_i = (x_i, y_i)\}_{i=1}^n$ is

$$\chi(\mathbf{c}, r) = \sum_{i=1}^n \left((\mathbf{c} - \mathbf{p}_i)^T (\mathbf{c} - \mathbf{p}_i) - r^2 \right)^2. \quad (6)$$

We wish to find r which minimizes χ so we solve

$$(\partial\chi/\partial r)(\mathbf{c}, r_0) = 0$$

for r_0 , obtaining

$$r_0^2 = \frac{1}{n} \sum_{i=1}^n (\mathbf{c} - \mathbf{p}_i)^T (\mathbf{c} - \mathbf{p}_i). \quad (7)$$

Then, find the minimum over \mathbf{c} by taking the partial with respect to \mathbf{c} , but after substituting $r = r_0$,

$$\begin{aligned} & \frac{\partial \chi}{\partial \mathbf{c}}(\mathbf{c}, r_0) \\ &= \frac{\partial}{\partial \mathbf{c}} \sum_{i=1}^n \left((\mathbf{c} - \mathbf{p}_i)^T (\mathbf{c} - \mathbf{p}_i) - \frac{1}{n} \sum_{j=1}^n (\mathbf{c} - \mathbf{p}_j)^T (\mathbf{c} - \mathbf{p}_j) \right)^2 \\ &= \frac{1}{n^2} \frac{\partial}{\partial \mathbf{c}} \sum_{i=1}^n \left(\sum_{j=1}^n (2\mathbf{c} - \mathbf{p}_i - \mathbf{p}_j)^T (\mathbf{p}_j - \mathbf{p}_i) \right)^2 \\ &= \frac{4}{n^2} \left[2 \left(\sum_{i,j,k} (\mathbf{p}_k - \mathbf{p}_i)(\mathbf{p}_j - \mathbf{p}_i)^T \right) \mathbf{c} + \sum_{i,j,k} (\mathbf{p}_i^T \mathbf{p}_i + \mathbf{p}_j^T \mathbf{p}_j)(\mathbf{p}_k - \mathbf{p}_i) \right]. \end{aligned}$$

Thus, the solution to $(\partial \chi / \partial \mathbf{c})(\mathbf{c}_0, r_0) = 0$ is

$$\begin{aligned} \mathbf{c}_0 &= -\frac{1}{2} \mathbf{A}^{-1} \mathbf{b} \text{ where } \mathbf{b} = \sum_{i,j,k} (\mathbf{p}_i^T \mathbf{p}_i - \mathbf{p}_j^T \mathbf{p}_j)(\mathbf{p}_k - \mathbf{p}_i) \text{ and} \\ & \mathbf{A} = \sum_{i,j,k} (\mathbf{p}_k - \mathbf{p}_i)(\mathbf{p}_j - \mathbf{p}_i)^T. \end{aligned} \quad (8)$$

Note that simplifying \mathbf{A} and \mathbf{b} yields single as opposed to triple sums, thereby reducing computational cost. Hence, the best fitting circle, using the metric defined in (6), has radius r_0 (7) (after substituting $\mathbf{c} = \mathbf{c}_0$) and is centered at \mathbf{c}_0 (8).

3.1 Calibration Algorithm

We describe in this section a complete method of calibrating a parabolic catadioptric sensor. We show that there are explicit formulas for calculation of the mirror center, focal length and aspect ratio. The only assumption of the data is that the images of at least three lines have been obtained.

Step 1. Obtain points. Start by obtaining images of points lying on lines in space. Assume the measured points are grouped into sets of points:

$$\mathbf{P} = \left\{ \left\{ \mathbf{q}_j^i = \begin{pmatrix} u_j^i \\ v_j^i \end{pmatrix} \right\}_{j=1}^{m_i} \right\}_{i=1}^n,$$

where for each $i = 1, \dots, n$ the set

$$\left\{ \mathbf{q}_j^i \right\}_{j=1}^{m_i}$$

is a collection of measured points lying on the same line image. If aspect ratio and skew are unknown, then for all i , $m_i \geq 5$; if aspect ratio and skew are known then for all i , $m_i \geq 3$. In all cases, $n \geq 3$. We discuss ways to obtain these points in Section 4 and conditions on the positions of the lines after the algorithm.

Step 2. Estimate aspect ratio and skew. There exists a single affine transformation (unique up to rotation and scale) which transforms a set of ellipses whose axes are parallel and aspect ratios identical into a set of corresponding circles. To find such a transformation we minimize the sum of residuals of circles fitted to the ellipses over the set of possible affine transformations. In the case where there is known to be no skew, then, we only need to find the aspect ratio. A transformation

$$\mathbf{T}(\alpha) = \begin{pmatrix} \alpha & 0 \\ 0 & \alpha^{-1} \end{pmatrix}$$

induces an aspect ratio of α^2 , the inverse of this transformation is just $\mathbf{T}(\alpha^{-1})$. The form of $T(\alpha)$ is chosen so that $\det T(\alpha) = 1$.

For a given aspect ratio α^2 , the centers and radii of the best fitting circles may be computed using the substitution $\mathbf{p}_j^i(\alpha) = \mathbf{T}(\alpha) \mathbf{q}_j^i$, yielding $\mathbf{c}_0^i(\alpha)$ and $r_0^i(\alpha)$. Let




Fig. 2. Intersection of three spheres constructed from line images yields a point on the mirror's axis a distance $2f$ above the image plane.

$$\chi(\alpha) = \sum_{i=1}^n \chi(\mathbf{c}_0^i(\alpha), r_0^i(\alpha)),$$

which is the sum of residuals over all circles as a function of the aspect ratio, and find α_0 such that $\chi(\alpha_0)$ is minimum over all α . We find $\chi(\alpha)$ by substituting $r_0^i(\alpha)$, then $\mathbf{c}_0^i(\alpha)$,

$$\begin{aligned} \chi(\alpha) &= \sum_{i=1}^n \frac{1}{m_i^2} \sum_{j=1}^{m_i} \left(\sum_{k=1}^{m_i} (2\mathbf{c}_0^i(\alpha) - \mathbf{p}_j^i(\alpha) - \mathbf{p}_k^i(\alpha))^T (\mathbf{p}_k^i(\alpha) - \mathbf{p}_j^i(\alpha)) \right)^2 \\ &= \frac{1}{\alpha^4} A(\mathbf{P}) + B(\mathbf{P}) + C(\mathbf{P})\alpha^4, \end{aligned}$$

after a lengthy derivation, where $A(\cdot)$, $B(\cdot)$, and $C(\cdot)$ are scalar functions of the sets of points \mathbf{P} (which can be explicitly obtained). The equation $(\partial \chi / \partial \alpha)(\alpha_0) = 0$ is linear in $\alpha^{1/8}$ and, thus, α_0 can be found such that $\chi(\alpha_0)$ is a minimum. Then, the aspect ratio is $a = \alpha_0^2$.

If the skew is unknown then we may use

$$\mathbf{T}(\alpha, \beta) = \begin{bmatrix} \alpha & \beta \\ 0 & \alpha^{-1} \end{bmatrix}.$$

Unfortunately the resulting $\chi(\alpha, \beta)$ and its partial derivatives yield high degree polynomials, but, nevertheless, can be minimized using gradient descent.

Once α_0 (and possibly β_0) has been found, apply the transformation $\mathbf{T}(\alpha_0)$ (or $\mathbf{T}(\alpha_0, \beta_0)$) to the measured points, i.e.,

$$\mathbf{p}_j^i = \begin{pmatrix} u_j^i \\ v_j^i \end{pmatrix} = \mathbf{T}(\alpha_0) \begin{pmatrix} u_j^i \\ v_j^i \end{pmatrix}.$$

Step 3 Fit circles to points. Using the algorithm described at the beginning of this section, compute the best fitting circles to each set $\{\mathbf{p}_j^i\}$ of image points. Let \mathbf{c}_i and r_i , respectively, be the center and radius of the i th circle.

Step 4. Find mirror center and focal length. Let ξ be the mirror center and f be the focal length. We know that a sphere constructed at each line image contains the point $(\xi_x, \xi_y, 2f)^T$, so if we intersect three such spheres, we obtain this point. For this reason, it is necessary to have at least three line images, see Fig. 2. We wish then to intersect all of these spheres constructed from line images. However, in the presence of noise the intersections of more than three of these spheres in space will almost certainly be empty. Hence, instead of finding the intersection, we find a point in space which minimizes the distance to

all of the spheres. We will actually minimize the sum of the squares of the difference of squared distances from the centers because this is more convenient for the same reasons we do so when fitting a circle to a set of points. It is easily verified that the sum of these squared distances is,

$$\chi(\xi, f) = \sum_{i=1}^n \left((\xi - \mathbf{c}_i)^T (\xi - \mathbf{c}_i) + 4f^2 - r_i^2 \right)^2.$$

This formula is suspiciously similar to our χ in (6). Using almost exactly the same analysis, we find that

$$f_0^2 = \frac{1}{4n} \sum_{i=1}^n \left(r_i^2 - (\xi - \mathbf{c}_i)^T (\xi - \mathbf{c}_i) \right)$$

and

$$\xi_0 = -\frac{1}{2} \mathbf{A}^{-1} \mathbf{b} \text{ with } \mathbf{b} = \sum_{i,j,k} \left(\mathbf{c}_i^T \mathbf{c}_j - r_i^2 - \mathbf{c}_j^T \mathbf{c}_k + r_j^2 \right) (\mathbf{c}_k - \mathbf{c}_i) \text{ and}$$

$$\mathbf{A} = \sum_{i,j,k} (\mathbf{c}_k - \mathbf{c}_i) (\mathbf{c}_j - \mathbf{c}_i)^T,$$

minimizes $\chi(\xi, f)$.

We have thus estimated the mirror center, the focal length, aspect ratio, and optionally skew. With the exception of the estimation of skew, this algorithm can be implemented using explicit formulas, no iterative minimization methods are required.

Before continuing, we examine the conditions under which we can calibrate from three lines (in the noiseless setting). To estimate the intrinsics, the intersection of three spheres constructed from line images must be two points. This does not hold when the circles' centers lie on a line and intersect in two points in the plane, i.e., they are coaxial. The two points of intersection in the plane, rotated in space about the line through the centers, is a circle equal to the intersection of all three spheres, not the desired two points. However, if the circles are not coaxial then the constructed spheres intersect in two points. That the line images are not coaxial is equivalent to the condition that the equivalent perspective line images do not all intersect in a single point.

3.2 Rectification

In this section, it is supposed that we have taken an image of a planar surface and we wish to reproject to an image plane parallel to this surface, i.e., perform a rectification. In order to do so, we must have calibrated the device and we must determine the orientation of this plane. We assume that on the planar surface are sets of parallel lines. Each set must contain at least two lines. There must be at least two sets in which the lines of one set are not parallel to the lines of the other set. The following algorithm describes a method of estimating the orientation of the plane from the horizon of the plane. The horizon is found from the vanishing points of the parallel lines.

This algorithm is applied in Section 4 and throughout the explanation of the algorithm we reference Fig. 5 to illustrate the steps of the algorithm.

Step 1. Obtain points and calibrate. At least two such sets must be obtained and each set must contain at least two line images. Assume at least three image points are measured from each line (five when either aspect ratio or skew are unknown). We group the points as in:

$$\left\{ \mathbf{P}_i = \left\{ \mathbf{l}_{i,j} = \left\{ \mathbf{p}_k^{i,j} \right\}_{k=1}^{m_{i,j}} \right\}_{j=1}^{n_i} \right\}_{i=1}^N,$$

where \mathbf{P}_i is a set of line images all of which are images of lines which are parallel to each other, and $\mathbf{l}_{i,j}$ is a set of points of some line image. Calibrate using the union of all the sets \mathbf{P}_i of line images. Let ξ and f be the mirror center and focal length

respectively. We assume from now on that all image points have been translated so that the mirror center is the origin.

See Fig. 5b. We show an example of an image in which pixels have been chosen lying on two sets of lines which are images of two sets of parallel lines in space. Circles are fitted to the pixels lying on the same line and calibration is performed using these circles.

Step 2. Fit line images to the point sets. The circles fitted to the points on the same line images before calibration are not necessarily line images satisfying the invariant (3). Find the best fitting line image to the set of points. We substitute $r^2 = \mathbf{c}^T \mathbf{c} + 4f^2$ into χ from (6), for general points $\{\mathbf{p}_i\}_{i=1}^n$,

$$\chi(\mathbf{c}) = \sum_{i=1}^n \left((\mathbf{c} - \mathbf{p}_i)^T (\mathbf{c} - \mathbf{p}_i) - \mathbf{c}^T \mathbf{c} - 4f^2 \right)^2. \quad (9)$$

The solution, again from a similar analysis, is

$$\mathbf{c}_0 = \frac{1}{2} \mathbf{A}^{-1} \mathbf{b} \text{ where } \mathbf{b} = \sum_{i=1}^n (\mathbf{p}_i^T \mathbf{p}_i - 4f^2) \mathbf{p}_i \text{ and } \mathbf{A} = \sum_{i=1}^n \mathbf{p}_i \mathbf{p}_i^T.$$

Fit a line image to each set of points $\{\mathbf{p}_k^{i,j}\}_{k=1}^{m_{i,j}}$. Let $\mathbf{c}_{i,j}$ be the center of the j th circle from the i th set.

See Fig. 5c. Line images are fitted to the points using the intrinsics calculated from calibration. Notice that their centers lie on two lines and that each set intersects in two points; they are therefore coaxial.

Step 3. Find the vanishing points. In order to estimate the horizon of the plane, which will give us its orientation, we need to find vanishing points on the horizon. In a parabolic projection, a set of parallel lines is projected to a system of coaxial circles (their centers lie on a line) which have two intersection points that are the pair of vanishing points. In a calibrated system, the axis of the coaxial system uniquely determines these two points; the intersection of any two line images whose centers lie on this line yields the vanishing points. The first step then is to determine this axis, we accomplish this by fitting a line to the centers of the line images in a parallel set. The mean of the centers lies on this line and the eigenvectors of the covariance matrix determine the direction. So, if

$$\bar{\mathbf{c}}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} \mathbf{c}_{i,j} \text{ and } \mathbf{C}_i = \sum_{j=1}^{n_i} (\mathbf{c}_{i,j} - \bar{\mathbf{c}}_i) (\mathbf{c}_{i,j} - \bar{\mathbf{c}}_i)^T,$$

and we let \mathbf{d}_i be the normalized eigenvector corresponding to the minimum eigenvalue of \mathbf{C}_i , then the implicit equation of the line is

$$(x \ y) \mathbf{d}_i - \bar{\mathbf{c}}_i^T \mathbf{d}_i = 0.$$

We choose the line image whose center is closest to the mirror center, the vanishing points are at the intersection of this line image and the perpendicular through its center, hence

$$\mathbf{v}_i^\pm = \left(\mathbf{d}_i^T \bar{\mathbf{c}}_i \pm \sqrt{(\mathbf{d}_i^T \bar{\mathbf{c}}_i)^2 + 4f^2} \right) \mathbf{d}_i$$

are the vanishing points ($\mathbf{d}_i^T \bar{\mathbf{c}}_i$ is the distance of the circle to the mirror center and the square root is the radius of the circle using the invariant (3)).

See Fig. 5d. The vanishing points are the intersection of the circles. Estimate the vanishing points by finding the line image whose center is on the line through the centers and closest to the image center. The vanishing points then lie on the intersection of this circle with the line perpendicular to the line through the centers and through the image center.

Step 4. Estimating the horizon. To find the horizon, which will give us the orientation of the plane parallel to all of the lines, we fit a line image (a circle satisfying the invariant equation) to the set of vanishing points. Use the same procedure described in Step 2 of this algorithm for fitting a line image to a set of points in a

calibrated system. Let \mathbf{h} be the center of the line image which is the horizon, and let r be its radius, then from the coefficients in (5), the orientation of the plane is

$$\hat{\mathbf{n}} = \left(-\frac{h_x}{r}, -\frac{h_y}{r}, \frac{2f}{r} \right).$$

This plane is parallel to all of the lines.

See Fig. 5e. A line image is fitted to the vanishing points; this circle is the horizon of the plane of the ceiling. The figure shows how one can estimate the azimuth and altitude of the plane from the center of the circle.

Again, this algorithm can be implemented without the use of iterative minimization methods and has completely explicit formulas. In particular, the eigenvectors of a 2×2 matrix are easily obtained, all other computations involve addition or inversion of a 2×2 matrix.

4 EXPERIMENTS

We have performed a number of experiments to study the robustness of the calibration algorithms developed in this paper. In the first experiment, we have run simulations to see what effect noise, field of view, and number of lines have on the accuracy of the calibration algorithm. Second, we compare our algorithm with a mirror-boundary calibration algorithm. Lastly, we demonstrate rectification of a scene with natural lines.

4.1 Simulations

In the following experiments, we generate a random scene consisting of measurements from artificially generated lines and artificially generated intrinsic parameters. We generate the point measurements in the following manner.

First, in a single run in an experiment, the image center is uniformly chosen from the square $[-240, 240] \times [-240, 240]$ and the focal length is constant at 120 pixels having by default an 180° field-of-view; this it to model a typical system which uses a CCD chip which has a dimension of 640×480 pixels, within which we can fit an 180° field-of-view mirror with an 120 pixel focal length.

Second, we generate a fixed number (20 by default) of line samples (not pixels yet): A plane is chosen with unit normal randomly and uniformly distributed on the sphere. This plane represents the projection of the line in space. If the projection of this line is within the selected field of view of the sensor, then, the line is accepted, otherwise it is rejected and another line sample is chosen until it is within the field of view. Next, an angle is randomly chosen to be the maximum angle subtended by any two rays sampled from the line. The angle is chosen from a distribution similar to a histogram of the maximum subtended angle of a pair of rays (calibrated image points) from the ceiling image (Fig. 5a) and from the calibration target images (Fig. 4). It is a normal distribution centered at 90° with standard deviation of 28° and truncated at 45° (an angle below which line images have too short of an arclength and which is less than those found in the real images) and 180° (the maximum possible angle subtended by two points on a line).

Third, from every line sample we choose a fixed number (20 by default) of point samples: A ray lying in the plane of the line is picked randomly and uniformly from within the field of view. This is the central ray about which points will be sampled on the line image. Rays are picked to have a uniform distribution of angle with the central ray up to half the maximum subtending angle chosen above; if a sampled ray does not lie within the field of view it is rejected and another is chosen. The rays are then projected to image coordinates using formula (1) and transformed using (2) with the chosen intrinsic parameters. Gaussian noise is added to pixel measurements; though this is not necessarily the best noise model, we are not at the moment aware of a more appropriate one.

Fig. 3e shows an example of a simulated scene with lines; in this case ten lines each with twenty points are randomly generated and perturbed with Gaussian noise having standard deviation of one.

Experiment 1. We wish in this experiment to determine the effect that pixel noise has on calibration. We create artificial image points as described above where the aspect ratio is 1, skew is 0 and both are assumed to be known. An independent, zero-mean, zero-correlated, two-dimensional, normal distribution of variable standard deviation is added to each pixel measurement. Then, we apply the calibration algorithm, using the assumed unit aspect ratio and skew (thus, without estimating aspect ratio or skew), to the simulated scene. The known parameters are compared with the parameters determined using the algorithm; an RMS error is computed over multiple runs of the experiment.

Fig. 3a and 3b shows the results of this experiment for a varying number of lines. Each curve shows the RMS error in the calibration results over 1,000 experiments as a function of noise; the number labeling each curve indicates how many lines were used to perform the calibration. We vary the standard deviation of the Gaussian noise in increments of $1/2$ pixel starting at $1/2$ pixel.

As we would expect, a greater amount of noise decreases the accuracy of the measurements. The question then is, by increasing the number of lines used to calibrate, can we increase the accuracy? For it may be easier to increase the number of lines used than to decrease the noise in our measurements. The fact that the curves with greater number of lines lie below those with fewer indicates that a greater number of lines does increase accuracy. Accuracy of focal length estimation is not increased greatly by additional lines. The error in image center, however, is much more affected by the number of lines. Consider that the image center is contained in a two-dimensional disk constructed from every line image in the image plane. It may therefore be estimated by intersecting all disks constructed from line images, and as the number of line images increases, the region of intersection becomes smaller, lowering the uncertainty of the location of the image center; the sphere intersection algorithm appears to share this property.

Experiment 2. In the final simulation, we show the effects that the field of view has on calibration. Though field of view is not something we may adjust once we have a sensor, it may be useful to know how it affects the performance. Again, we create artificial image points as described previously. We assume that the skew is 0, and thus we do not estimate it.

Fig. 3c and 3d shows the results of calibration while varying the field of view in increments of 45° starting at 90° up to 270° . The two curves show the RMS error in the calibration results over 1,000 experiments as a function of the field of view. The dashed curve shows the errors when the aspect ratio is estimated prior to image center and focal length calibration. The solid curve shows the errors when only image center and focal length are calibrated. In the case where the aspect ratio is estimated, points have been transformed beforehand so that the aspect ratio is 1.1.

Interestingly, when aspect ratio is known, we see that variation in the field of view has relatively little effect on the accuracy of image center estimation, i.e., the image center error curves are relatively flat. The image center can be localized from the intersection of two line images which happen to be lines, so the closer a line image is to a line, the better a localization we can obtain of the image center. By decreasing the field of view, we do not exclude these lines and, so, the image center estimation is not affected by variation in the field of view.

However, the lines closer to the fronto-parallel horizon are excluded with low field of view. These happen to be the lines which localize the focal length well, for the height of the spheres they generate determines the focal length. For this reason, as the field of view increases, the accuracy of focal length estimates increases.




Fig. 3. (a) and (b) RMS errors for image center and focal length under varying noise. Different curves are for different numbers of lines. In the graph of focal length error, the curves are too close to label; a curve lying above another has a greater number of lines. (c) and (d) RMS errors under varying field of view. Full lines are for known aspect ratio, dashed lines unknown aspect ratio. (e) An example of a simulated scene with 10 lines, 20 points each, field of view 180° . Gaussian noise of 1 pixel standard deviation added.

There is negligible difference in error of the estimation of focal length between the known and unknown aspect ratio cases. This could be because the spheres constructed using the estimated aspect ratio will have little variation in their total diameter (giving the focal length). However, for the image center there is much greater error because error in calibrating the aspect ratio—which is directly related to the the average arclength of the line images—causes greater error in extrapolating the trajectories of the line images beyond the support of the measurements (i.e., with greater arclength it is easier to estimate the circle that an arc is part of), greatly

affecting the positioning of the sphere and as a consequence having a greater effect on the estimation of the image center.

4.2 Comparison

In this section, we compare the performance of our algorithm with a mirror boundary algorithm which comes with the software delivered with the mirror. We will examine the figures listed in Table 1. Six images were taken using a a folded catadioptric camera prototype (S360m) made by RemoteReality which has a resolution of 640×480 and a field of view of 180° . Fig. 4 shows one of the six images; each image is of a calibration target viewed from different




Fig. 4. (a) One of six images taken with the S360m. (b) Virtually projected grid superimposed onto image using two calibrated parameters, our algorithm (black points) and the mirror boundary algorithm (white points). The circles are centered at the calibrated image centers and have radius twice the calibrated focal length, colors correspond as above.

orientations. In each image, the points on the grid have been manually obtained and put into correspondence with six virtual lattices. From the points on the lattice we can construct all possible sets of points which lie on the same line on the grid. We do so in the vertical, horizontal and diagonal directions for each grid. For each image this gives us a set of points lying on the same line.

We apply the mirror boundary algorithm three times to each of the six frames (for a total of 18 measurements). The first row of the first column shows the averaged results of the calibration with standard deviation; the next row shows the average error in fitting line images (using (9) but divided by the number of points and summed over all line images and divided by the number of line images) to the lines on the grid. This is a measure of how well the projection, with intrinsics found using the mirror boundary calibration algorithm, projects lines in space to circles satisfying the invariant (3). The third row will be explained at the end of this section.

Then, we apply the sphere intersection algorithm (where we assume the aspect ratio is 1 and skew is 0) but *not* to the lines on the grid. We first locate points by applying a Laplacian filter to the image, threshold, and then group into blobs. The positions of the blobs are then given to a random sample consensus [25] algorithm to find circles from random samples of triples. From each image the 20 circles with the greatest support are taken. Of these 120, (20 from the six frames), a random subset of 50 are chosen and used to calibrate the sensor. This entire procedure is performed 50 times and the results are shown with standard deviation in the first row of the second column of Table 1. The second column shows the RMS error of fitting line images to the lines of the grid which is less than the same RMS error for the mirror boundary algorithm.

Because of lack of ground-truth and in order to enable a comparison between the two methods we devised an additional algorithm. This evaluation algorithm minimizes using gradient descent the point distance between the image points and the projection of points of a virtual grid obtained from the lattice. The minimization is performed over the six virtual grids' pose orientation and the intrinsics. The intrinsics which minimize the residual defined in this way are shown in the first row of the third column. The second row shows the average line residual using these intrinsics. The third row shows the average point distance over all the points of all six grids; this is the quantity which was minimized.

This evaluation metric searches for the intrinsic and pose parameters which project points of a virtual grid closest to image points that are projections of the points of a real grid in space. The

intrinsics so found are not a ground truth, however what are found are values which model the sensor with the least amount of error. It therefore gives us intrinsic parameters against which we may compare the intrinsic parameters obtained from both the mirror boundary algorithm and the sphere intersection algorithm. In the image plane, the image center calculated using sphere intersection is closer to the one obtained by the evaluation algorithm. The difference in focal lengths is also less for the sphere intersection algorithm.

We end with one more comparison. We take the evaluation algorithm above and we perform the same minimization of point distances except that we do not minimize over the intrinsics, just over the six grids' pose orientations. We hold the intrinsics constant at the values obtained from the two algorithms. We therefore have a measure of how well the two intrinsics model the sensor. In the third row of the first column, we show the minimum average error in modeling the sensor using the intrinsic parameters obtained from the mirror boundary algorithm. In the third row of the second column, we show the minimum average error in modeling the sensor using the intrinsic parameters obtained from the sphere intersection algorithm. The RMS of projected point error is less using the intrinsics calculated from the intersection of spheres than from the mirror boundary.

According to the two metrics, RMS of fitted line images and RMS of point projection error of the virtual grid, the algorithm introduced in this paper models this device with less error than a mirror boundary algorithm. Also, the intrinsic parameters obtained with the evaluation algorithm are closer to that obtained by the sphere intersection algorithm.

4.3 Rectification

In this last experiment, we demonstrate rectification of a "natural" image—not a calibration target—a picture taken of the ceiling of our laboratory. On the ceiling of our lab is a steel structure for holding wires, etc. which has parallel metal bars in two directions. We wish to reproject the image onto a virtual image plane which is parallel to the surface of the ceiling. Fig. 5a shows the original image and Fig. 5f shows the rectified image. Figs. 5b, 5c, 5d, and 5e illustrate the steps of the rectification algorithm. Notice that in the rectified image, even though we did not use the assumption that the bars of the grid were perpendicular (which they actually are) the lines of the grid in the reprojection are at 90° angles.




Fig. 5. (a) Original image of ceiling, (b) circles fitted to seven line images and used to calibrate, (c) line images fitted to pixels and lines fit to circles' centers, (d) vanishing points calculated; (e) horizon determined from four vanishing points and orientation (azimuth and elevation) estimated from the position of the horizon's center, and (f) rectified image produced by projecting to the plane parallel to the ceiling.

TABLE 1
Results of Applying the Different Algorithms to the Six Images

Algorithm	Mirror Boundary (averaged)	Sphere intersection (averaged)	Projection error minimization
Image center (ξ) \pm std. dev.	(327.6, 236.7) \pm 3.5	(323.7, 238.2) \pm 3.0	(317.3, 238.1)
Focal length (f) \pm std. dev.	123.0 \pm 2.2	115.7 \pm 0.5	116.1
RMS of geometric distance of points to fitted line images	0.519	0.383	0.3980
RMS of projected point error	2.791	2.200	2.193

See text for details. All figures are in units of pixels.

5 CONCLUSION

In this paper, we have presented novel geometric properties of the projection induced by an optical system consisting of a parabolic mirror and an orthographic camera. We have shown that lines in three-dimensional space, including those lines at infinity, project to a particular family of circles. An invariant of this family of circles, a quantity changed only by the focal length and the mirror center, enables us to estimate these intrinsic parameters from a subset of the family. We have provided a novel single algorithm which can calibrate all of the intrinsic parameters of a parabolic projection, including aspect ratio and skew, from at least three input lines in general position. The algorithm requires absolutely no metric information. In addition, we present a rectification algorithm to estimate the unknown normal of a plane from one uncalibrated view using only the image of the line at infinity. It is known that no rectification of a plane is possible from one uncalibrated view taken by a conventional camera without metric information.

We have tested each of the algorithms and in particular compared the calibration to another algorithm based on the image of the mirror's boundary. We show that it yields an estimate of the intrinsic parameters which models the true projection induced by the sensor quantitatively better.

In our ongoing work, we study the geometric properties of all catadioptric systems with unique viewpoints, our first results have appeared in [7].

ACKNOWLEDGMENTS

This work has been supported by the US Army Research Office/Multidisciplinary University Research Initiative DAAH04-96-1-00007, US National Science Foundation ISS-0083209, US National Science Foundation GER93-55018, US National Science Foundation CDS-97-03220, Defense Advanced Research Projects Agency-ITO-MARS-130-1303-4-534328, and Advanced Network.

REFERENCES

- [1] E. Hecht and A. Zajac, *Optics*. 3rd ed., Addison-Wesley, 1997.
- [2] G. Toomer, "Diocles On Burning Mirrors," *Sources in the History of Mathematics and the Physical Sciences*, Springer-Verlag, 1976.
- [3] S. Nayar, "Catadioptric Omnidirectional Camera," *IEEE Conf. Computer Vision and Pattern Recognition*, pp. 482-488, June 1997.
- [4] S. Baker and S. Nayar, "A Theory of Single-Viewpoint Catadioptric Image Formation," *Int'l J. Computer Vision*, vol. 35, pp. 175-196, 1999.
- [5] T. Svoboda, T. Pajdla, and V. Hlavac, "Epipolar Geometry for Panoramic Cameras," *Proc. 5th European Conf. Computer Vision*, pp. 218-231, 1998.
- [6] A. Bruckstein and T. Richardson, "Omniview Cameras with Curved Surface Mirrors," *Proc. IEEE Workshop Omnidirectional Vision*, pp. 79-86, June 2000, Originally published as Bell Labs technical memo, 1996.
- [7] C. Geyer and K. Daniilidis, "A Unifying Theory for Central Panoramic Systems," *Proc. Sixth European Conf. Computer Vision*, pp. 445-462, 2000.

- [8] C. Taylor, "Video Plus," *Proc. IEEE Workshop Omnidirectional Vision*, pp. 3-10, June 2000.
- [9] P. Sturm, "A Method for 3D-Reconstruction of Piecewise Planar Objects from Single Panoramic Images," *Proc. IEEE Workshop Omnidirectional Vision*, pp. 119-126, June 2000.
- [10] T. Boulton, "Remote Reality Demonstration," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 966-967, June 1998.
- [11] Y. Onoe, K. Yamazawa, H. Takemura, and N. Yokoya, "Telepresence by Real-Time View-Dependent Image Generation from Omnidirectional Video Streams," *Computer Vision and Image Understanding* vol. 71, pp. 588-592, 1998.
- [12] N. Winters, J. Gaspar, G. Lacey, and J. Santos-Victor, "Omnidirectional Vision for Navigation," *Proc. IEEE Workshop Omnidirectional Vision*, pp. 21-28, June 2000.
- [13] A. Leonardis and M. Jogan, "Robust Localization Using Eigenspace of Spinning-Images," *Proc. IEEE Workshop Omnidirectional Vision*, pp. 37-46, June 2000.
- [14] R. Benosman, E. Deforas, and J. Devars, "A New Catadioptric Sensor for Panoramic Vision of Mobile Robots," *Proc. IEEE Workshop Omnidirectional Vision*, pp. 112-118, June 2000.
- [15] V. Nalwa, "A True Omnidirectional Viewer," technical report, Bell Labs, Holmdel, NJ 1996.
- [16] Y. Yagi, "Omnidirectional Sensing and Its Application," *IEICE Trans. Information & Systems*, vol. 3, pp. 568-579, 1999.
- [17] *Proc. IEEE Workshop Omnidirectional Vision*. K. Daniilidis, ed., June 2000.
- [18] R. Benosman and S. Kang, *Panoramic Vision*. Springer-Verlag, 2000.
- [19] Y. Yagi, S. Kawato, and S. Tsuji, "Real-Time Omnidirectional Image Sensor (copis) for Vision-Guided Navigation," *Trans. Robotics and Automation*, vol. 10, pp. 11-22, 1994.
- [20] S. Kang, "Catadioptric self-calibration," *IEEE Conf. Computer Vision and Pattern Recognition*, pp. 1-201-207, June 2000.
- [21] S. Nene and S. Nayar, "Stereo with Mirrors," *Proc. Int'l Conf. Computer Vision*, pp. 1087-1094, Jan. 1998.
- [22] I. Coope, "Circle Fitting by Linear and Nonlinear Least Squares," *J. Optimetric Theory Applications*, vol. 76, pp. 381-388, 1993.
- [23] L. Moura and R. Kitney, "A Direct Method for Least-Squares Circle Fitting," *Computer Physics Comm.*, vol. 64, pp. 57-63, 1991.
- [24] Z. Zhang, "Parameter-Estimation Techniques: A Tutorial with Application to Conic Fitting," *Image and Vision Computing*, vol. 15, pp. 59-76, 1997.
- [25] M. Fischler and R. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Comm. ACM*, vol. 24, pp. 381-395, 1981.

► For more information on this or any other computing topic, please visit ourDigital Library at <http://computer.org/publications/dlib>.