Calibration of Cameras with Radially Symmetric Distortion

Jean-Philippe Tardif *
DIRO, Université de Montréal
Canada

Peter Sturm †
INRIA Rhône-Alpes
38330 Montbonnot St Martin, France

Abstract

We present a theory and algorithms for plane-based calibration and pose recovery of general radially distorted cameras. By this we understand cameras that have a distortion center and an optical axis such that the projection rays of pixels lying on a circle centered on the distortion center, form a right cone centered on the optical axis. The camera is said to have a singular viewpoint (SVP) if all such view cones have the same vertex (the optical center), otherwise we speak of non-SVP, and each view cone may have its own optical center on the optical axis. This model encompasses the classical radial distortion model, fisheyes, most central or non-central catadioptric cameras, but also cameras with radially symmetric caustics.

Calibration consists in the estimation of the distortion center, the opening angles of all view cones and their optical center. We present two approaches of computing a full calibration from dense correspondences of a single or multiple planes with known euclidean structure. The first one is based on a geometric constraint linking view cones and associated ellipses in the calibration plane; calibration of the view cones can be solved by determining the closest point to a set of hyperbolas. The second approach uses existing plane-based calibration methods to directly calibrate individual view cones. A simple distortion correction algorithm for calibrated SVP images is given. Preliminary experiments show convincing results.

1. Introduction

In the last few years, we have seen an increasing interest in non-conventional cameras and projection models, going beyond affine or perspective projection. There exists a large diversity of camera models; many of them specific to certain types of projections [1, 13] or families of cameras such as central catadioptric systems [2, 8, 3, 6]. All these models are described by a few intrinsic parameters, much like the classical pinhole model, possibly enhanced with radial or decentering distortion coefficients. Calibration methods exist for all these models, and they are usually tailor-made for them, i.e. can not be used for any other projection model. Several works address the calibration problem from an opposite point of view, by adopting a very generic imaging model that incorporates most commonly used cameras [9, 5, 17, 10, 15]. In the most general case, cameras are modeled by attributing an individual projection ray to each individual pixel. Such a model is highly expressive, but it is difficult to obtain a stable calibration of cameras with it.

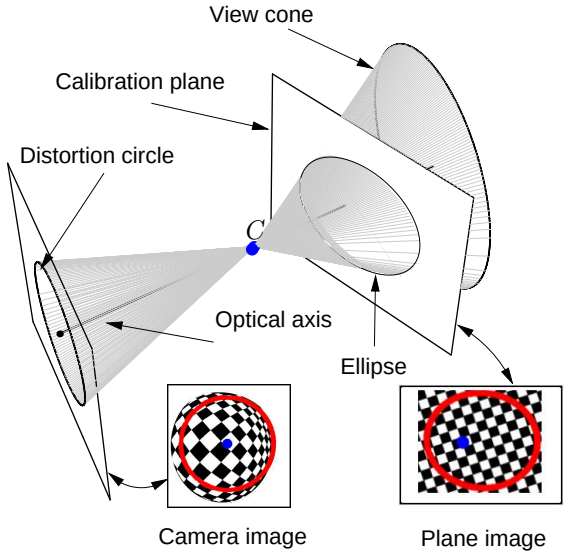

Figure 1: Our camera model (see text for explanations). The inlayed illustrations show the distortion center (in blue) and a distortion circle for a true image taken with a fisheye, and the corresponding calibration ellipse overlaid on a pattern shown on the calibration plane.

In this paper, we propose a simple camera model (and associated calibration methods) that hopefully offers a good compromise: it is sufficiently general to model many common camera types, but has much fewer parameters than the above generic model, thus making calibration much easier and more stable. We model cameras using the notions of a **distortion center** in the image and an **optical axis** in 3D. For cameras with **radially symmetric distortion**, the projection rays associated with pixels lying on a same circle centered on the distortion center, lie on a right **viewing cone** centered on the optical axis (cf. fig. 1). This encompasses many common camera models, such as pinhole (modulo aspect ratio and skew), the classical polynomial radial dis-

*tardifj@iro.umontreal.ca
†peter.sturm@inrialpes.fr

tortion model, fisheyes, or any catadioptric system whose mirror is a surface of revolution, and for which the optical axis of the perspective (or affine) camera looking at it is aligned with the mirror's revolution axis. Our model comprises **central** cameras (SVP), where all viewing cones have the same vertex (the **optical center**), but also **non-central** ones (NSVP), for which the viewing cones' vertices lie anywhere on the optical axis. In the latter case, we may speak of one optical center per viewing cone.

Problem statement. We want to calibrate cameras based on the above model, from one or several images of a calibration plane in unknown positions. The input to the calibration algorithms is a dense matching between the plane(s) and the camera image, and the euclidean structure of the plane(s). From this, we compute, for all viewing cones, their focal length (equivalent to the opening angle). Our algorithms assume a known position of the distortion center, but we also show how to estimate it, using repeated calibrations for different candidates. Calibration also comprises a pose estimation problem: estimating the orientation of the optical axis (relative to a calibration plane) and the location of each viewing cone's vertex on it.

Organization. A geometric study of our model is presented in §2, together with our first calibration approach. The second approach, based on the standard plane-based calibration method, is described in §3. In §4, we give an algorithm for performing perspective image rectification based on calibration results. Several practical issues and experimental results are presented in §5 and §6, respectively.

2. Geometry

2.1. One Distortion Circle

Let us consider one distortion circle in the image plane. We suppose that we have determined the ellipse on the calibration plane that is mapped to that circle via the camera's projection function (see §5). If we knew the position of the camera's optical center relative to the calibration plane, then we could compute the cone that has the optical center as vertex and that contains the above calibration ellipse. That cone has several interesting properties: its axis is the camera's optical axis and it is a right cone, i.e. rotationally symmetric with respect to its axis. From the cone, the focal length of the considered distortion circle can be easily computed (the cone's opening angle equals the field of view).

In practice, we do not know the optical center's position relative to the calibration plane. In the following, we show geometrical relations between the calibration ellipse, the optical center and the optical axis of the camera. When talking about optical center, we mean the optical center per distortion circle; they all lie on the optical axis and in the SVP case, they are identical.

Without loss of generality, we assume that the calibration plane is the plane $Z = 0$, and that the calibration ellipse is given by the matrix $\mathrm{diag}(a, b, -1)$, with $b \geq a > 0$, i.e. the X-axis is the ellipse's major axis. Our aim is to provide constraints on the position of the optical center, as well as on the orientation of the optical axis, from this ellipse.

Let us first state a well-known result. Consider a right cone (whose vertex is a point with real coordinates) and its intersection with a plane. For now, we only consider the case where the intersection is an ellipse (the case of the hyperbola will be discussed later). It is easy to prove that the orthogonal projection of the cone's vertex onto the plane, lies on the ellipse's major axis (cf. fig. 2a and §5). This implies that the cone's vertex lies in the plane that is orthogonal to the ellipse's supporting plane and that contains its major axis.

For our problem, this means that the optical center must lie in the plane $Y = 0$ (since the ellipse lies in plane $Z = 0$ and has the X-axis as major axis). We may further constrain its position $\mathbf{C} = (X, 0, Z, 1)^\mathsf{T}$, as follows [4]. The cone with $\mathbf{C}$ as vertex and that contains the calibration ellipse, is given by ($\propto$ represents equality up to scale):

$$\Lambda \propto \begin{pmatrix} a\,Z^2 & 0 & -a\,X\,Z & 0 \\ 0 & b\,Z^2 & 0 & 0 \\ -a\,X\,Z & 0 & a\,X^2 - 1 & Z \\ 0 & 0 & Z & -Z^2 \end{pmatrix}.$$

For this cone to be a right one, its upper left 3×3 matrix $\bar{\Lambda}$ must have a double eigenvalue. The three eigenvalues are:

$$b\,Z^2\,,\ \frac{a\,X^2 + a\,Z^2 - 1 \pm \sqrt{4a\,Z^2 + \left(-a\,X^2 - a\,Z^2 + 1\right)^2}}{2}$$

The second and third eigenvalues can not be equal for real values of X and Z (besides in the trivial case $X = Z = 0$). The first eigenvalue is equal to the third one if $Z = 0$ and to the second one if:

$$a\,b\,X^2 + b(a - b)\,Z^2 + (a - b) = 0. \qquad (1)$$

This equation tells us that the optical center lies on a conic in the plane $Y = 0$, given by the following matrix:

$$\Psi = \begin{pmatrix} a\,b & & \\ & b(a-b) & \\ & & a - b \end{pmatrix}.$$

This is a hyperbola, since $(a - b) < 0$. Furthermore, its asymptotes correspond to the direction of the two cylinders that contain the calibration ellipse.

Let us now consider the orientation of the optical axis. Due to (1), let us consider an optical center with:

$$Z = \pm\sqrt{\frac{a\,b\,X^2 + a - b}{b(b - a)}}.$$

In proceedings of *The 6th the Sixth Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras, (OMNIVIS'2005)*, in conjunction with ICCV'2005, october 2005.

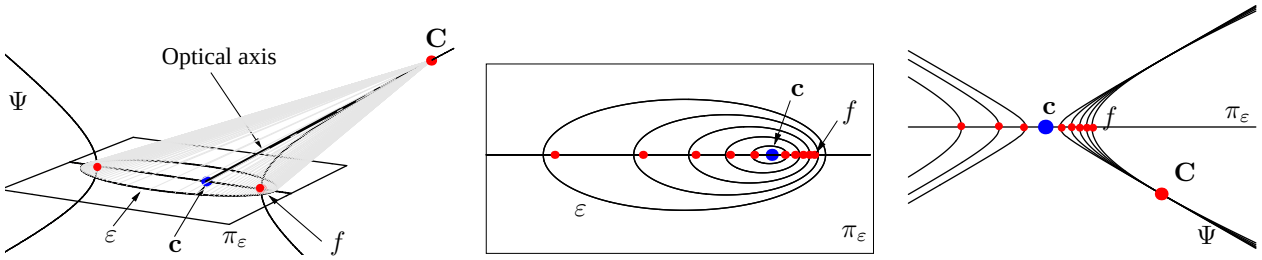

Figure 2: Illustrations of the geometry of viewing cones, calibration ellipses and location of optical center. cf. text. **a**) Complete illustration for one viewing cone. **b**) View of the calibration plane, showing many cones' calibration ellipses. Note that their major axes are collinear. **c**) Side view of the hyperbolas associated with many calibration ellipses.

Here, we exclude the case $a = b$, which would correspond to the camera looking straight at the calibration plane.
The direction of the cone's axis is given by the eigenvector associated with the single eigenvalue of $\bar{\Lambda}$, augmented with a homogeneous coordinate 0:

$$\begin{pmatrix} \pm\sqrt{b(b-a)(a\,b\,X^2+a-b)} \\ 0 \\ a\,b\,X \\ 0 \end{pmatrix}. \qquad (2)$$

We now show that the cone's axis is identical with the tangent of the hyperbola Ψ in the optical center $\mathbf{C}$, which is given by (in the plane $Y = 0$):

$$\Psi \begin{pmatrix} X \\ Z \\ 1 \end{pmatrix} = \begin{pmatrix} a\,b\,X \\ \mp\sqrt{b(b-a)(a\,b\,X^2+a-b)} \\ a-b \end{pmatrix}.$$

Its point at infinity is (still in the plane $Y = 0$):

$$\begin{pmatrix} \pm\sqrt{b(b-a)(a\,b\,X^2+a-b)} \\ a\,b\,X \\ 0 \end{pmatrix}$$

i.e. it is identical with the point given in (2). Hence, for an optical center on Ψ, the optical axis is directly given by the associated tangent.
The case where the intersection between a cone and the calibration plane yields a hyperbola, given by $\mathrm{diag}(a, -b, -1)$, can be dealt with in a similar fashion. This typically occurs with very wide angle cameras. Once again, the calibration hyperbolas have their major axes aligned together. We can show that the possible viewpoints lie on an ellipse given by $\mathrm{diag}(a\,b, b(a+b), -a-b)$ and that the optical axis is tangent to it. For simplicity's sake, the rest of the article concentrates on the elliptic case; nevertheless, everything holds when some intersections are hyperbolas.

2.2. Multiple Distortion Circles

So far, we have shown that for an individual distortion circle, the associated viewing cone can be determined from the associated calibration ellipse, up to 1 degree of freedom (location on the hyperbola Ψ). We now show how to get a unique solution, when considering several distortion circles simultaneously. Let us first note that calibration ellipses corresponding to different distortion circles, are not independent: their major axes are collinear (cf. fig. 2b) [1]. Their centers are not identical however (unless they are all circles, i.e. if the camera looks straight at the plane).
Let Ψ_d be the hyperbolas for different distortion circles, given in the same coordinate frame. In the case of a single viewpoint camera, the optical center must lie on all these hyperbolas. Furthermore, the optical axis is tangent to all of them. This implies that all hyperbolas touch each other (have a double intersection point) in the optical center. This is illustrated in figure 2c. A naïve algorithm would compute the hyperbolas for all ellipses and seek their single intersection/contact point. The drawback of this situation is that very little noise can cause two hyperbolas to have no real intersection point at all, instead of a double one.
Consider now the NSVP case: to each distortion circle and viewing cone, corresponds a separate optical center. Hence, the hyperbolas won't have a single contact point anymore. However, the optical axis is shared by all viewing cones. Hence, it is the single (in general) line that is tangent to all hyperbolas. Furthermore, the individual optical centers are its contact points with the associated hyperbolas.

2.3. Calibration and Pose Estimation

A simple calibration method consists in computing the 3D point which is closest in average to all hyperbolas (see next paragraph). This gives the camera's optical center (relative to the calibration plane). Then, viewing cones can be spanned with individual calibration ellipses, and the focal

[1] This constraint is non-linear, but can be enforced when fitting the ellipses in cases where the correspondences with the calibration plane have large errors, or not uniformly distributed around the curve. It is not shown here due to lack of space.

In proceedings of *The 6th the Sixth Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras, (OMNIVIS'2005)*, in conjunction with ICCV'2005, october 2005.

lengths for all distortion circles are computed based on their opening angles. In the NSVP case, we would first compute the optical axis. A plausible criterion would be to find the line $\mathbf{L}$ that minimizes the sum of squared distances between itself and the closest tangent line to each hyperbola parallel to $\mathbf{L}$. In the following, this calibration approach is referred as the Right Cone Constraint method (RCC).

Closest point to the hyperbolas. Computing the orthogonal distance of a point to a general conic requires solving fourth degree polynomial [19]. Using this to compute the closest point to our set of hyperbolas is not very practical. Instead, we minimize a simpler cost function: the closest point $\mathbf{q}$ is found by solving:

$$\min_{\mathbf{q},\mathbf{q}_d} \sum_d dist(\mathbf{q},\mathbf{q}_d)^2, \text{ subject to } \mathbf{q}_d^\mathsf{T}\,\Psi_d\,\mathbf{q}_d = 0,$$

i.e. we also estimate one point per hyperbola Ψ_d that will, after convergence, be the orthogonal projection of $\mathbf{q}$ on Ψ_d. The problem is solved using the *Minimize* function of *Mathematica*. Since the function and constraints are polynomial, it uses a cylindrical algebraic decomposition that guarantees a global minimum [18].

3. Calibration with the IAC

The RCC approach relied on pose recovery from the image of one plane to calibrate. In practice, if many calibration planes are available, one would want to use them to increase robustness. We present another approach that first computes the calibration (from one or many images of planes) and then recovers the pose. The approach uses well-known results on plane-based calibration for perspective cameras [20, 16]. Indeed, it is possible to see the viewing cones in terms of many *perspective cameras*, with different focal lengths but identical principal point. In the SVP case, their extrinsic parameters are also identical, whereas an NSVP camera can be modeled by adding a translation along the optical axis per viewing cone.

Calibration. Let us consider the distortion circle of radius d and one image of a calibration plane. From point correspondences between pixels on this circle and points on the calibration plane (on the calibration ellipse), we can compute a plane-to-image homography H_d. For simplicity, let us assume that image coordinates have been translated to put the distortion center at the origin. The homography can then be decomposed such that:

$$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \propto \mathsf{H}_d \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \mathsf{K}_d\,\mathsf{R} \begin{pmatrix} 1 & 0 & \\ 0 & 1 & -\mathbf{t} + t_d\mathbf{r}_3 \\ 0 & 0 & \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \tag{3}$$

where (x, y) is a calibration point, (u, v) a pixel on the distortion circle, and R and $\mathbf{t}$ a rotation matrix and translation vector representing camera pose (same for all d). The scalar

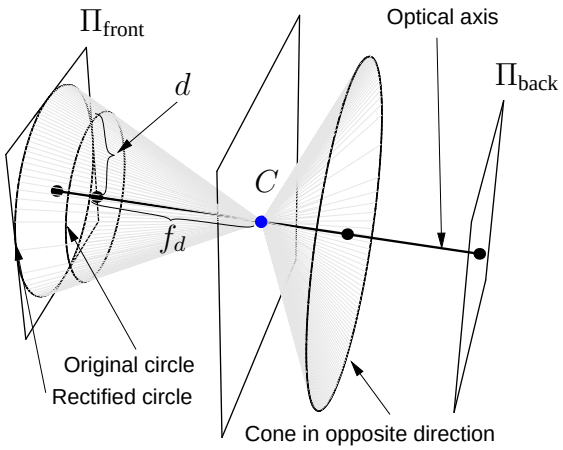

Figure 3: Viewing cones can also be seen as individual perspective cameras with different focal length. A rectified image can be obtained by projecting the distortion circles (which lie in different planes) on one plane Π_{front} (or Π_{back} for a field of view larger than $180°$).

t_d allows to model translational displacement of individual viewing cones along the optical axis (given by $\mathbf{r}_3^\mathsf{T}$, the third row of R), which is needed for NSVP cameras. For SVP cameras, we set $t_d = 0$ for all d. As for K_d, it is a calibration matrix[2] $\text{diag}(f_d, f_d, 1)$, where f_d is the focal length associated with the considered distortion circle. We may interpret the relation between d and f_d as a distortion function applied to a perspective camera whose undistorted image plane is π_{front} (cf. fig. 3).

Note that this parameterization only accounts for viewing cones with field of view smaller than $180°$. Larger fields of view can be modeled by adding a reflection to the rotational component of the pose, $\mathsf{R}' = \text{diag}(1, 1, -1)\mathsf{R}$, and a corresponding image plane π_{back}.

From H_d, we first compute K_d, using the approach of [20, 16]. Of course, this can be done using the homographies given for multiple images of the calibration plane.

Once the calibration is known for each viewing cone, the pose R and $\mathbf{t}$ can be computed from the homography of any distortion circle, using [14]. In the SVP case, the pose is the same for all d, and we may "average" the different estimates or better, non-linearly optimize the pose and calibrations simultaneously for all d. In the NSVP case, we first compute R, which is the same for all d. As for the position of optical centers ($\mathbf{t}$ and the t_d), we must fix one t_d, e.g. $t_0 = 0$.

[2] As mentioned in the introduction, this model does not include a skew between pixel axes or an aspect ratio different from 1. Also, it assumes that the distortion center is at the principal point. Generalizing this would be straightforward though.

In proceedings of *The 6th the Sixth Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras, (OMNIVIS'2005)*, in conjunction with ICCV'2005, october 2005.

Then, from H_0, we can compute $\mathbf{t}$ and finally from the H_d, the scalars t_d. This time, non-linear optimization is recommended and straightforward to perform.

3.1. Computing the Distortion Center

Until now, we have assumed, for both algorithms, that the distortion center was known; this information was used to select the distortion circles. Tests with noiseless simulated data showed that the calibration may be quite sensitive to a bad choice of distortion center; as for real cameras, using the image center as an approximation was not satisfying in general. Hence, the distortion center should be estimated as part of the calibration process. The sensitivity of calibration we observed in simulation suggests that it should be possible to estimate the distortion center rather reliably, which was confirmed in practice.

We used the following heuristics to define an optimization criterion for the distortion center. Let us apply the IAC approach of §3 with several images as input. The plane-based calibration for each distortion circle is then capable of estimating a principal point, besides the focal length f_d. It seems plausible that the better the assumed distortion center was, the closer the estimated principal points will be to it. Since plane-based calibration is applied on images centered on the assumed distortion center, we can consider the average norm of the estimated principal points (on per distortion circle) as a measure for the goodness of the center.

Figure 4 shows the values of this measure, computed for distortion center positions on a 60×60 grid around the image center, for real cameras. The shape of the cost surface indicates that we can find the optimum distortion center using a simple steepest descent type method. We implemented such an approach that accurately finds the distortion center within a couple of minutes of computation. Note that the second row of figure 4b shows that, although the principal points used to plot it were computed individually per distortion circle, they are very densely clustered (average distance to assumed distortion center of less than 3 pixels). This suggests a high stability of the calibration.

4. Image Rectification

Once the calibration of an SVP camera is known, an image can be perspectively rectified. Then, straight lines in the scene appear straight in the image.

Rectification is done by placing a virtual perspective camera at the actual camera's optical center. Let K^v and $\mathsf{R}^v_{3\times3}$ represent the virtual camera's calibration and orientation. A naïve approach for image creation is to render each pixel of the original (distorted) image into the virtual (distortion-free) image, and then fill out the holes by interpolation (cf. fig. 3). It is well known that an inverse approach is better. We achieve this by inverting the relation between d and the view angle $\theta(d)$. As seen in figure 11, this function is generally simple (close to linear), so easily invertable (see §4). Let $D(\theta)$ denote this inverse function. One pixel $\mathbf{q}^v$ in the rectified image is backprojected in space with $q^w = (\mathsf{K}^v \mathsf{R}^v)^{-1} q^v$. Then, we compute the angle θ between this pixel and the Z-axis (the optical axis of the original camera). Finally, the corresponding position in the original image is given by $\mathsf{K}_{D(\theta)}\, \mathbf{q}^w$.

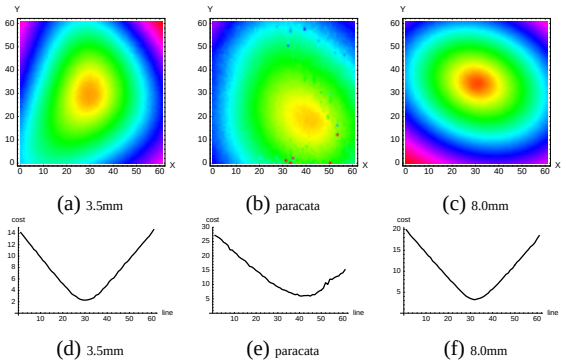

Figure 4: Plots of the goodness measure for the distortion center, obtained for three tested lens (cf. §6). **a,b,c**) 60×60 grid around the image center (yellow meaning smaller). **d,e,f**) One slice of the grid, through the best position.

5. Practical Issues

5.1. Dense Camera–Plane Correspondences

The easiest approach we found to get dense correspondences between the calibration plane and the camera is to use a flat screen as plane. We used a simple coded structured light algorithm [12], which consists in successively displaying patterns of horizontal and vertical black and white stripes on the screen to encode the position of each screen pixel (cf. fig. 5). Then, for each camera pixel, we identify the corresponding position on the calibration plane by decoding the observed intensities in each pattern. When performed in a controlled environment (low-constant ambient lighting, screen of high contrast and resolution), the accuracy of such a method is reasonably good (around ± 2 pixel of error on average). Since the points located on the distortion circles are given in floating point coordinates, we compute their correspondences by a weighted sum of the correspondences recovered for the four closest image pixels.

5.2. Omnidirectional Cameras

There are several issues worth mentioning for omnidirectional cameras. If the field of view is larger than $180°$, then some distortion circles will have viewing cones that actually approach planes. Their pose can not be estimated the same way as for true cones. These can be detected as the ones

In proceedings of *The 6th the Sixth Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras, (OMNIVIS'2005)*, in conjunction with ICCV'2005, october 2005.

(a) (b) (c)

Figure 5: Projected patterns for correspondences are horizontal and vertical black and white stripes. Images taken with **a**) the Goyo 3.5 mm, **b**) catadioptric, and **c**) paracatadioptric camera.

whose correspondences on the calibration plane are close to collinear. They can then be discarded from the actual calibration procedure and we may attribute them an unknown focal length. In the case of the IAC, rank deficient homographies are discarded resulting once again in an unknown focal length. Their actual value can be determine afterward by interpolation (see §6)

In the NSVP case, we still need to compute the offset t_d on the optical axis of such (almost) "viewing planes", since it may differ from that of other viewing cones. This is simple once the optical axis has been computed using other distortion circles and their exact opening angle has been determined: the cones' offset can be computed such that they go through the extracted correspondences in the calibration plane.

6. Experiments

We used a wide-angle Goyo 3.5 mm lens combined with a CCTV A201bc Basler camera, a Cosmicar 8.0 mm with little distortion, a paracatadioptric camera built with a Cosmicar 12.5 mm (referred to as "paracata" in the text), and also a homemade catadioptric device built from a Fujinon 12.5 mm lens, pointed at a roughly spherical mirror (cf. fig. 7). The calibration plane of known euclidean structure was a 21 inch CRT screen in all cases, except for the paracatadioptric camera where a multimedia projector was carefully placed in a fronto-parallel position w.r.t. to a wall. Even though the alignment was not perfect and the camera self-occluded, it did not affecte the solution significantly. The only non-linear optimizations that were performed to obtain the following results are in the hyperbola intersection algorithm and the pose estimation for the IAC approach.

Figure 10 gives the computed focal length of the 3.5 mm, 8.0 mm and paracatadioptric lenses, w.r.t. the distance d to the distortion center, using both methods. The wide-angle camera could already be calibrated from a single image of the screen with both approaches (cf. fig. 9a,b for the RCC), although better results were obtained using five images and the IAC approach. The paracatadioptric camera was calibrated with the two approaches with very similar results (cf. fig. 9c); however, the RCC algorithm gave the best results.

Figure 7: Catadioptric camera built from a Basler A201bc camera with a Fujinon 12.5 mm lens pointed at a roughly spherical mirror.

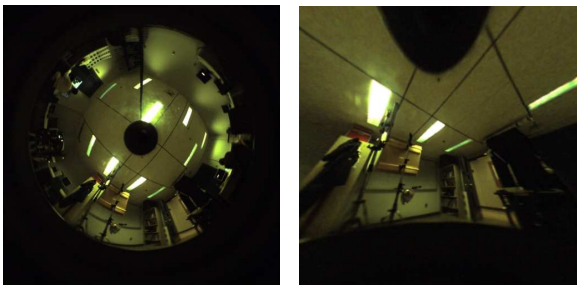

(a) (b)

Figure 8: Image rectification of the paracatadioptric camera. **a**) Original image. **b**) Rectified image for a rotated camera.

Indeed the radius for which the focal length is 0 corresponds better to the measurement of the correspondences' colinearity. The radial configuration of the catadioptric camera was not perfect. Nevertheless, the distortion center could be found and a satisfying calibration could be obtained with both methods. The IAC approach gave the best results because it could take advantage of up to eight images, which is more robust to the imperfect configuration of the camera. We also observed that the calibration is more stable for the lens with the wider field of view. Indeed, when there is very little distortion in the image, the hyperbolas' curvatures are similar, which induces more instability for the recovered camera pose. We also calibrated the 8.0 mm with the OpenCV library [11], and found the recovered focal lengths to be inside the result's uncertainty interval. Image rectification also yielded almost identical results.

In practice, only a subset of distortion circles are used for calibration; others can then be extrapolated or interpolated from a polynomial fitting of the data. Let us define this polynomial p; from the camera model, it is best to ensure that its derivative at 0 (corresponding to the distortion center) is 0. This constraint is due to the symmetry of the distortion model, and can be enforced by using a polynomial of the form $p(d) = a_0 + a_1 d^2 + ... + a_{n-1} d^n$. In practice, polynomials of degree 3 appeared to be sufficient. To handle the case of omnidirectional cameras more appropriately, the interpolation is carried out with the view angle instead of the focal length. In this case, a polynomial passing through 0 can also be fitted (see fig. 11).

In proceedings of *The 6th the Sixth Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras, (OMNIVIS'2005)*, in conjunction with ICCV'2005, october 2005.

(a) (b) (c)

Figure 6: Image rectification. **a**) Original images. **b**) Rectified image for the Goyo 3.5 mm. **c**) Rectified image for the home-made catadioptric camera. Small inset images show rectification of the border regions.

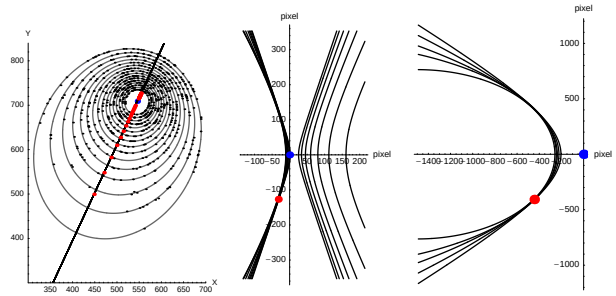

(a) (b) (c)

Figure 9: Calibration with the RCC approach. **a**) Fitted ellipses for the Goyo lens and **b**) corresponding hyperbolas and computed intersection. **c**) For the paracatadioptric camera, the intersection between the calibration plane and the cones yielded ellipses and hyperbolas, constraining the viewpoint to lie respectively on hyperbolas and ellipses.

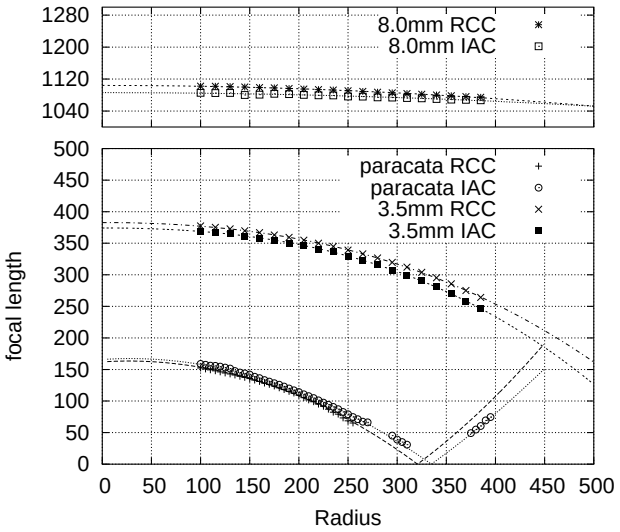

Figure 10: Recovered focal length (in pixel) from the two algorithms and extrapolated values from polynomial fitting of the data for the tested cameras.

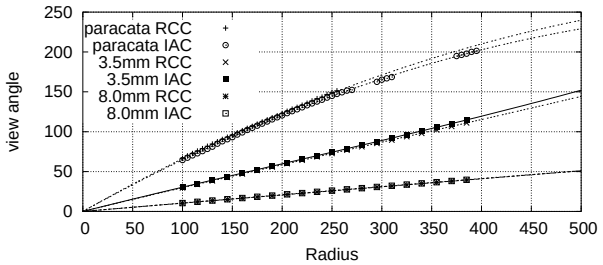

Figure 11: View angle in degrees for the tested cameras.

Evaluating the results based on the reprojection error can lead to biased conclusions in the case of a generic model. Indeed, the model offers more freedom which allows to fit the data better. This analysis goes together with the comparison between SVP and NSVP constraints and the displacement t_d of the viewpoints on the optical axis. This topic is to be explored more thoroughly, but the preliminary results obtained with the IAC approach indicate that our model is useful (see table 2). They show that the paracatadioptric and to a lesser extent the 3.5 mm are probably NSVP. The displacement along the optical axis confirms this observation; the shape of the curves also leads us to believe that it is not a result of overfitted data (see fig. 12).

More meaningful quantitative results were obtained for the

	Position			Angle		
Algorithms	p_{01}	p_{12}	p_{02}	a_{01}	a_{12}	a_{02}
Ground truth	5	5	10	0°	0°	0°
RCC	4.99	5.09	10.08	2.25°	0.89°	2.91°
IAC	4.89	5.13	9.99	0.6°	0.4°	1.1°

Table 1: Result for pose estimation. The camera was moved to three positions with known relative motion. Coefficients p_{ij} and a_{ij} denote the distance (in centimeters) and relative angle (in degrees) between camera positions i and j.

Camera	Constraint	
Algorithms	SVP	NSVP
paracata	9.10	1.01
3.5 mm	5.18	2.23
8.0 mm	2.20	1.35

Table 2: Comparison of the average reprojection error for different constraint of the viewpoint.

Goyo lens, using a pose estimation procedure. Using a translation stage, the camera was moved to three positions with known relative motion (no rotation, known translation). Using the calibration information (obtained using other images), the pose of the camera relative to the calibration plane was computed for all three positions (the NSVP configuration being very similar in all three cases). From this, the relative motions were computed and compared to the ground truth. The results presented in table 1 show a good stability for both methods.

Finally, images from the three panoramic cameras were rectified based on the calibration results (cf. fig. 6 and 8). Especially for the wide-angle Goyo lens (with little NSVP), the rectification seems to be very good, even towards the image borders (cf. the inset image in fig. 6b). As for our paracatadioptric camera, the rectification is very good, although not perfect, a likely result of its NSVP. Finally, the rectification of our home-made catadioptric device is also surprisingly good for a large part of the image, especially around the borders. The remaining distortions in the center were found to be caused by a small bump on the "mirror's" surface.

7. Summary and Conclusion

We have proposed new calibration approaches for a camera model that may be a good compromise between flexibility and stability for many camera types, especially wide-angle ones. The RCC approach is theoretically very interesting but its practical usability remains limited. This is due to the fact that only one calibration plane can be used directly. We also intend to perform a better analysis of its stability in the future.

The IAC algorithm, especially when used with many images, showed greater stability. It is also the basis for the distortion center estimation which is an important issue of the calibration.

Our approach may be especially suitable for unknown configurations (SVP/NSVP, mirror equation) or slightly misaligned catadioptric systems.

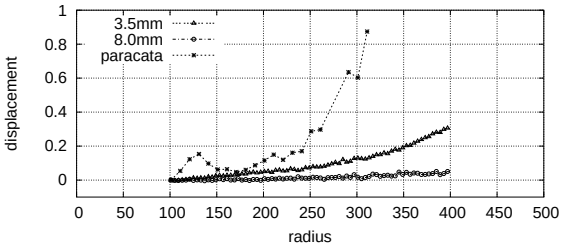

Figure 12: Displacement (in mm) along the optical axis for the 8.0 mm, the 3.5 mm and the paracatadioptric. The general curves' form of the two last leads us to think that the NSVP optimization is not a result of overfitting.

References

[1] R. Bajcsy, S.-S. Lin. True single view point cone mirror omnidirectional catadioptric system. ICCV 2001. 1

[2] S. Baker, S.K. Nayar. A Theory of Single-Viewpoint Catadioptric Image Formation. IJCV, 35(2), 1–22, 1999. 1

[3] J.P. Barreto, K. Daniilidis. Unifying image plane liftings for central catadioptric and dioptric cameras. OMNIVIS 2004. 1

[4] W. Boehm, H. Prautzsch. *Geometric Concepts for Geometric Design*. A.K. Peters, 1994. 2

[5] G. Champleboux, S. Lavallée, P. Sautot, P. Cinquin. Accurate Calibration of Cameras and Range Imaging Sensors: the NPBS Method. ICRA 1992. 1

[6] D. Claus, A.W. Fitzgibbon. A Rational Function Model for Fish-eye Lens Distortion. CVPR 2005. 1

[7] W. Gander, G.H. Golub, R. Strebel. Fitting of Circles and Ellipses. *BIT*, 34, 556–577, 1994.

[8] C. Geyer, K. Daniilidis. A Unifying Theory for Central Panoramic Systems and Practical Applications. ECCV 2000. 1

[9] K.D. Gremban, C.E. Thorpe, T. Kanade. Geometric Camera Calibration using Systems of Linear Equations. ICRA 1988. 1

[10] M. Grossberg, S. Nayar. A general imaging model and a method for fi nding its parameters. ICCV 2001. 1

[11] Intel Open Source Computer Vision Library. URL http://www.intel.com/research/mrl/research/opencv/. 6

[12] J. Salvi, J. Pagès, J. Batlle. Pattern codifi cation strategies in structured light systems. *Pattern Recognition*, 37(4), 827–849, 2004. 5

[13] D.E. Stevenson, M.M. Fleck. Nonparametric correction of distortion. TR 95–07, University of Iowa, 1995. 1

[14] P. Sturm. Algorithms for plane-based pose estimation. CVPR 2000. 4

[15] P. Sturm, S. Ramalingam. A generic concept for camera calibration. ECCV 2004. 1

[16] P. Sturm, S. Maybank. On Plane-Based Camera Calibration. CVPR 1999. 4

[17] R. Swaminathan, M. Grossberg, S. Nayar. Caustics of catadioptric cameras. ICCV 2001. 1

[18] E.W. Weisstein et al. Cylindrical algebraic decomposition. http://mathworld.wolfram.com/CylindricalAlgebraicDecomposition.html. 4

[19] Z. Zhang. Parameter estimation techniques: A tutorial with application to conic fitting. TR 2676, INRIA, 1995. 4

[20] Z. Zhang. A Flexible New Technique for Camera Calibration. PAMI, 22(11), 1330–1334, 2000.

4