Multi-view geometry of 1D radial cameras and its application to omnidirectional camera calibration

SriRam Thirthala Marc Pollefeys
Dept of Computer Science
UNC-Chapel Hill
$\{tvn, marc\}$@$cs.unc.edu$

Abstract

We study the multi-view geometry of 1D radial cameras. A broad a class of both central and non-central cameras, such as fish-eye and catadioptric cameras, can be reduced to 1D radial cameras under the assumption of known center of radial distortion. For cameras in general configuration, we introduce a quadrifocal tensor that can be computed linearly from 15 or more features seen in four views. From this tensor a metric reconstruction of the 1D cameras as well as the observed features can be obtained. In a second phase this reconstruction can then be used as a calibration object to estimate a non-parametric non-central model for the cameras. We study some degenerate cases, including pure rotation. In the case of a purely rotating camera we obtain a trifocal tensor that can be estimated linearly from 7 points in three views. This allows us to obtain a metric reconstruction of the plane at infinity. Next, we use the plane at infinity as a calibration device to non-parametrically estimate the radial distortion. We demonstrate the results of our approach on real and synthetic images.

1. Introduction

There has been a growing interest in the vision community for omnidirectional cameras. By providing a wider field of view these cameras are better suited for a range of applications such as visualization and robot navigation [4]. Two different approaches are used to obtain single images with a very wide field of view. The first approach consists of using a fish-eye lens which typically yields a field of view of around 180 degrees. The second approach consists of using a camera in combination with a curved mirror to obtain a cata-dioptric camera. The mirror surface is typically a quadric that is symmetric around the optical axis of the camera. When the center of projection of the camera is placed at one of the foci, an image with a single center of projection is obtained [2]. Otherwise, the catadioptric sensor corresponds to a non-central camera [23, 14]. The main topic of this paper is the calibration of omnidirectional cameras.

Many different techniques have been proposed to calibrate omnidirectional cameras. Most approaches are specific to one particular type of camera. For *fish-eye lenses* most approaches that have been proposed require some knowledge of the scene [21, 24, 6, 3, 16]. An approach based on pure rotation was proposed in [29]. By using a stratified approach in which the calibration model is gradually refined Pajdla et al. [18] have recently proposed a fully automatic approach that does not require any specific knowledge of the scene structure or the camera motion. For *cata-dioptric* cameras different types of techniques have also been proposed. Some techniques require knowledge about the scene [9, 30], others require knowledge about the camera motion [15]. Geyer and Daniilidis [8] proposed an approach that estimates the camera motion as well as the intrinsics of a parabolic cata-dioptric system. There has also been work on calibrating non-central cata-dioptric cameras [1, 18]. Besides this, there has also been work on generalized camera models. All the approaches that have been proposed for this require knowledge about the scene structure or the camera motion [10, 22].

In this paper we present two approaches for calibration of omnidirectional cameras based on multi-linear constraints between projection of point features. The first approach requires 15 or more point correspondences in four views and can deal with central and non-central cameras. The second approach works only for central cameras that perform pure rotation, but only requires 7 points in three views and is therefore ideally suited for automated calibration. Both approaches use non-parametric models of distortion so that all types of radially symmetric cameras can be modeled. The method requires to know the center of radial distortion. Both approaches consist of two main steps. First, the scene is reconstructed independently of the unknown radial distortion. Then, the reconstructed scene is used as a calibration object to estimate a non-parametric model of ra-

dial distortion. A different distortion can be estimated for each image so that different cameras can be calibrated simultaneously or the camera intrinsics can vary.

2. Radial 1D Camera

Suppose that the center of radial distortion is known. For most omnidirectional cameras the image center is a good approximation for the center of distortion [28]. [1] The image can then be transformed such that the center of radial distortion is the origin.

Consider a point in the world $\mathbf{X}$ that projects onto $\mathbf{x}_d = (x_d, y_d, 1)^T$ in the distorted (input) image. Further let $\mathbf{C}$ be the camera center. Because of large unknown and possibly varying distortion, the point $\mathbf{X}$ does not lie on the ray passing through $\mathbf{C}$ and $\mathbf{x}_d$ (see Figure 1(a)).

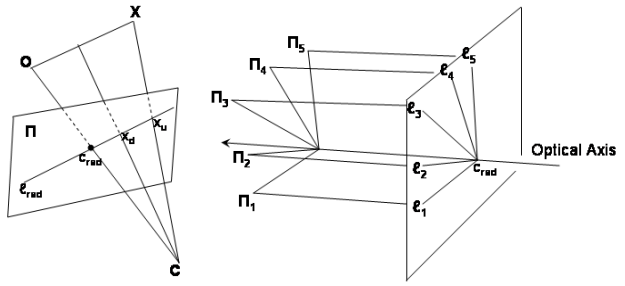

(a) (b)

Figure 1. *Radial 1D Camera*

However, consider the line passing through the center of radial distortion and $\mathbf{x}_d$ in the image ($\mathbf{l}_{rad} = \mathbf{x}_d \times \mathbf{c}_{rad}$). The undistorted image point (one that would have been obtained if the camera had followed a pin-hole projection model) $\mathbf{x}_u$, would lie on this line. This is because though the *distance* of an image point from the center of radial distortion is not preserved by radial distortion, the *direction* (which is what the radial line $\mathbf{l}_{rad}$ encodes) is. If instead of back-projecting a ray, we back-project the line $\mathbf{l}_{rad}$ using the camera center, it would contain the ray passing through $\mathbf{C}$ and $\mathbf{x}_u$, and thus would contain $\mathbf{X}$.

Thus, by representing the distorted image as a 1D image of radial lines passing through the center of radial distortion, we can factor out the *unknown deviation* from the pinhole model (which is along the radial line), but preserve the *known information* (direction of radial line). The radial 1D camera can be thought of as projecting the bundle of planes containing the optical axis onto the bundle of lines passing through the $\mathbf{c}_{rad}$ (Figure 1(b)). A radial line can be represented as $\mathbf{l} = (y, -x)^T$ if $\mathbf{c}_{rad}$ has been mapped to the origin. Note that a radial 1D camera can be obtained for most single effective viewpoint cameras (standard pin-hole cameras, low radial distortion cameras, fish-eye lenses, catadioptric cameras [2]). Infact we can deal with non-central cameras also. The only requirement is that all points that lie in one plane, of the bundle around the optical axis, project onto the same radial line (passing through $\mathbf{c}_{rad}$). For catadioptric systems, this corresponds to the requirement that all the normals on the mirror have to be contained within radial planes. This constraint is automatically satisfied for mirror shapes that are symmetric around the optical axis.

Definition: The radial 1D camera represents the mapping of a point in P^3 onto a radial line in the image. A $P^3 \rightarrow P^1$ projective mapping, it can be represented by a 2×4 matrix and has 7 degrees of freedom.

The projection of a 3D point $\mathbf{X}$ on a radial line $\mathbf{l}$ using radial camera $\mathbf{P}$ is then given by:

$$\lambda \mathbf{l} \quad = \quad \varepsilon \mathbf{P} \mathbf{X} \quad \text{with } \varepsilon = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \tag{1}$$

with λ a non-zero scale factor. Note that a point $\mathbf{O}$ on the optical axis does not have a proper image in P^1 as we obtain $\mathbf{PO} = (0, 0)^T$. The backprojection of a radial line to a plane is given by:

$$\lambda \mathbf{\Pi} \quad = \quad \mathbf{P}^T \mathbf{l} \tag{2}$$

3. Radial Quadrifocal Tensors

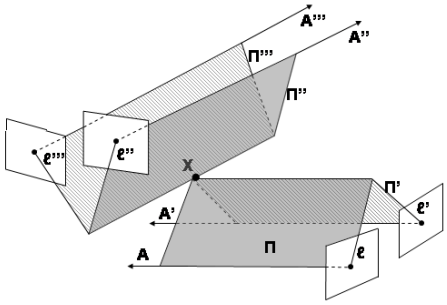

Figure 2. *The Quadrifocal Constraint*

Let us examine the possible multi-view constraints using these 1D radial cameras in general configuration. Note that we only have back-projected planes and no back-projected rays (as the distance information from the $\mathbf{c}_{rad}$ is unknown, only the radial line is preserved as a set). Three planes in 3D-space always intersect at a point and hence 3 views give us no constraints. However, four planes intersecting at a common point yields a non-trivial constraint (Figure 2).

[1] However if additional information is known, for example one can see the rim of a fish-eye lens or the rim of a curved mirror, the center of the rim can be used as a better approximation for the center of radial distortion.

Thus we have multi-view constraints among four 1D radial cameras.

Consider a point $\mathbf{X}$ in $\mathbf{P}^3$ that projects onto the radial lines, $\mathbf{l}, \mathbf{l}', \mathbf{l}'', \mathbf{l}'''$. Then the radial projection equations (Eq. 1), can be collected in the following matrix equation:

$$\begin{bmatrix} \varepsilon\mathbf{P} & \mathbf{l} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \varepsilon\mathbf{P}' & \mathbf{0} & \mathbf{l}' & \mathbf{0} & \mathbf{0} \\ \varepsilon\mathbf{P}'' & \mathbf{0} & \mathbf{0} & \mathbf{l}'' & \mathbf{0} \\ \varepsilon\mathbf{P}''' & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{l}''' \end{bmatrix} \begin{bmatrix} \mathbf{X} \\ -\lambda \\ -\lambda' \\ -\lambda'' \\ -\lambda''' \end{bmatrix} = \mathbf{0} \qquad (3)$$

Since we know that a solution exists, the right null-space of the 8×8 measurement matrix should have non-zero dimension, which implies that the determinant of the matrix has to be zero.

Following the approach of [26], expansion of the determinant yields the quadrilinear constraint for 1D radial cameras:

$$\mathbf{l}_i \mathbf{l}'_j \mathbf{l}''_k \mathbf{l}'''_l \mathbf{Q}^{ijkl} = 0 \qquad (4)$$

$\mathbf{Q}^{ijkl}$ is the $2 \times 2 \times 2 \times 2$ homogenous quadrifocal tensor of four 1D cameras. We use the Einstein summation convention in which indices repeated in covariant and contravariant positions denote implicit summations.

The radial quadrifocal tensor has $2 \times 2 \times 2 \times 2 - 1 = 15$ degrees of freedom. Subtracting from these the degrees of freedom required to describe four uncalibrated 1D radial cameras ($4 \times (2 \times 4 - 1) - (4 \times 4 - 1) = 13$), we observe that the radial quadrifocal tensor has *only 2* internal constraints. Compare this to $(80 - 29)$ internal constraints for the quadrifocal tensor of 4 perspective views.[2] The radial quadrifocal tensor can thus be linearly estimated given 15 corresponding quadruplets. Given more than 15 corresponding quadruplets, a linear least squares solution can be obtained. It can further be proven that there exist no higher-order tensors for 1D cameras.

3.1. 3D Reconstruction

3.1.1. Projective Reconstruction

We now consider the problem of 3D reconstruction of points whose correspondences have been specified across the input images. Given a radial quadrifocal tensor, we can easily compute the four uncalibrated camera matrices [12]. For every valid radial quadrifocal tensor, two non-equivalent projective reconstructions are obtained. As we can not disambiguate between them at this stage we will carry them through to the metric reconstruction stage and potentially the radial calibration where in general only a single solution will yield consistent results. Once the projection matrices have been recovered, points in 3D can be reconstructed by back-projecting planes. This corresponds to computing the right nullspace of the following matrix:

$$\begin{bmatrix} \mathbf{P}^T\mathbf{l} & \mathbf{P}'^T\mathbf{l}' & \mathbf{P}''^T\mathbf{l}'' & \mathbf{P}'''^T\mathbf{l}''' \end{bmatrix}^T \qquad (5)$$

Since only three planes are required to define a point uniquely in 3D space, we can infact reconstruct all points seen in atleast three views.

3.1.2. Metric Reconstruction

The dual absolute quadric, $\mathbf{\Omega}^*_\infty$ encodes both the absolute conic and the plane at infinity. To upgrade our reconstruction to metric, we need to estimate this degenerate quadric in the projective frame in which the cameras and the points have been determined [27, 19]. $\mathbf{\Omega}^*_\infty$ projects into the radial 1D image as,

$$\tilde{\mathbf{K}}\tilde{\mathbf{K}}^T = \tilde{\omega}^* = \mathbf{P}\mathbf{\Omega}^*_\infty\mathbf{P}^T \qquad (6)$$

with $\tilde{\mathbf{K}} = \begin{bmatrix} f_x & s \\ 0 & f_y \end{bmatrix}$ the upper 2×2 part of the calibration matrix. Using the assumptions of (i) *zero skew* ($s = 0$) (ii) *known aspect ratio* ($f_y = af_x$), we obtain 8 linear constraints on $\mathbf{\Omega}^*_\infty$, from the 4 views. Since $\mathbf{\Omega}^*_\infty$ is a 4×4 homogenous symmetric matrix it has 9 d.o.f (10 upto scale). Using the additional rank-3 constraint we obtain a fourth-degree equation ($det\mathbf{\Omega}^*_\infty = 0$) and thus obtain upto 4 solutions. Only positive semi-definite solutions for the absolute quadric have to be considered. If more than one solution persists, we can generate multiple alternative metric reconstructions and disambiguate them later by verifying the radial symmetry in the next section. If $\mathbf{\Omega}^*_\infty$ is decomposed as $\mathbf{\Omega}^*_\infty = \mathbf{H}diag(1, 1, 1, 0)\mathbf{H}^T$, then $\mathbf{H}^{-1}$ is the point homography that takes the projective frame to the metric frame [13].

3.2. Radial calibration

Once a metric reconstruction has been obtained using the 1D radial property of the camera, it can be used to calibrate the remaining unknowns of the projection. In this section we will present a non-parametric approach to calibrate central and non-central radially symmetric cameras. This process can be done independently for each image and it is thus possible to calibrate four different cameras -or a camera with different settings- using a single quadrifocal tensor.

For each cameras, all reconstructed feature points can be represented in a cylindrical coordinate system relative to the optical axis of the camera, i.e. (ρ, ϕ, z). The origin along the z-axis can in a first phase be chosen arbitrarily. Because we assume radial symmetry, the ϕ coordinate is irrelevant for us. The goal of our calibration procedure is to obtain an expression for rays, $\mathbf{r}$, in the ρz-plane as a function of the radius r, i.e. $\mathbf{r}(r) : a_1(r)\rho + a_2(r)z + a_3(r) = 0$. This can

[2] In fact, for perspective cameras, the radial quadrifocal tensor corresponds to the upper $2 \times 2 \times 2 \times 2$ part of the full quadrifocal tensor.

be done by fitting lines to all the points that have (almost) the same r value.

3.3. Synthetic Experiments

We will now describe simulations that we carried out to test the validity and robustness of reconstruction using the quadrifocal tensor. The following 4 cameras were chosen: a pin-hole camera looking at a spherical mirror, a pin-hole camera looking at a hyperbolic mirror (satisfying the single effective view-point condition [2]), a perspective camera and a fish-eye camera. Only the points which were imaged in all the 4 views were considered (a total of 2300 points were imaged, see Figure 4). To every point in every image, Gaussian noise with $\sigma = 1$ pixel (in an image of 2000x2000 pixels) was added. The final 4 images are shown in Figure 3. The 4 cameras were modeled to have zero skew and unit aspect ratio.

In Figure 5 the results of the metric reconstruction are shown. Only the difference vectors, between the ground truth and the reconstruction obtained, are plotted. The ratio of the RMS reconstruction error and the standard deviation of the ground truth point set is less than 1 percent. This ratio grew to around 3-5 percent when noise of $\sigma = 2$ pixel was introduced.

In the second phase, we perform radial calibration. For each camera, the first phase would have given us a precise optical axis in metric space. We select an arbitrary point on the optical axis and compute (ρ, z) pairs for each reconstructed point . In this 2D coordinate system, all points which project onto the same radial circle, in the input image, should lie on the same line. Given a sufficiently dense set of points, we can estimate these incoming rays (see Figure 3). Note that all the incoming rays for the three central cameras (views 2,3 and 4) pass through a point on the optical axis, as expected. This happens without enforcing any explicit constraint. For a non-central camera (view 1), the envelope of rays corresponds to the caustic of a spherical mirror as expected.

4. Radial Trifocal Tensor

Suppose that three optical axes, $\mathbf{A}, \mathbf{A}'$ and $\mathbf{A}''$ intersect at $\mathbf{C}$. Also, suppose that a 3D point $\mathbf{X}$ projects onto the lines $\mathbf{l}, \mathbf{l}'$and $\mathbf{l}''$ in the three views (Figure 6) . Consider the plane, $\mathbf{\Pi}$ containing $\mathbf{C}$ and the line $\mathbf{l}$ (corresponding to the back-projection of the radial line $\mathbf{l}$). Similarly, one has the planes $\mathbf{\Pi}^{'}$ and $\mathbf{\Pi}^{''}$. Note that for every 3D point, $\mathbf{X}$, the corresponding planes back-projected from the 3 views intersect in the line passing through $\mathbf{C}$ and $\mathbf{X}$. Three planes in 3D space intersecting in a line is a non-trivial constraint. This non-trivial constraint between the three 1D radial views is

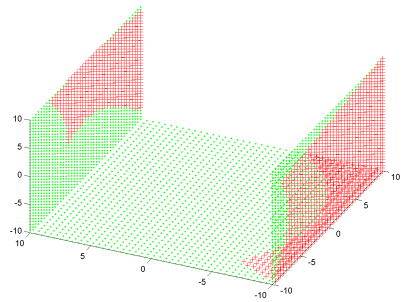

Figure 4. *The scene that is imaged by the 4 cameras. Plus (+) signs mark the points imaged in all cameras and dot (.) mark points which werent.*

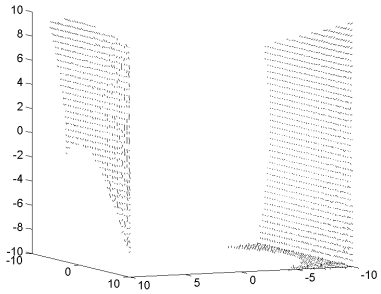

Figure 5. *These line segments connect the ground-truth points to the metric reconstruction obtained. Compare the length of the segments to the extent of the scene to get an idea of the error in reconstruction.*

encoded by the radial trifocal tensor. We can now formulate this constraint mathematically.

Without loss of generality we can assume that the three optical axes intersect in the origin $(0, 0, 0, 1)^T$. Since $\mathbf{PC} = (0, 0)^T$, the 1D radial cameras whose optical axes contain the origin must have the following form $\mathbf{P} = [\tilde{\mathbf{P}}|\mathbf{0}]$. Let $\tilde{\mathbf{X}}^T$ correspond to the first three coefficients of $\mathbf{X}$. In

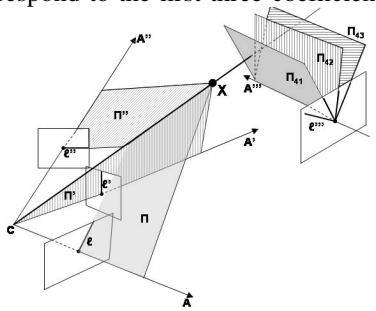

Figure 6. *The Trifocal Constraint*

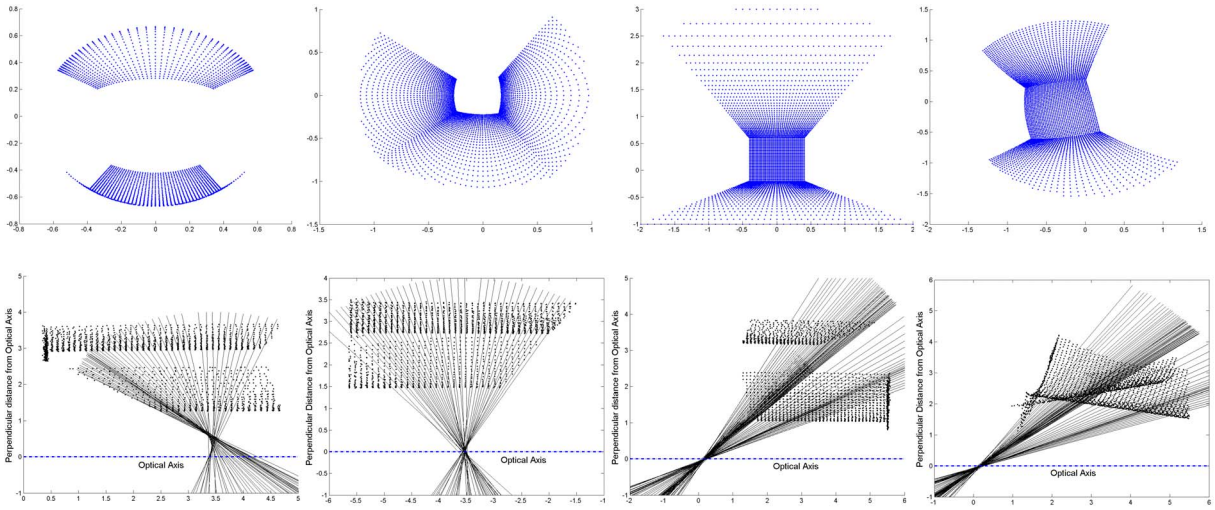

Figure 3. *(L to R) Spherical mirror, Hyperbolic mirror, Pin-Hole Camera, Fish-eye Lens.* ***Top****:Images obtained by the four cameras* ***Bottom****: (ρ, z) plots for reconstructed features and estimated incoming rays. Notice the caustic of the spherical-mirror camera (extreme left)*

this case, the first 6 rows of Eq. (3) can be rewritten as:

$$\begin{bmatrix} \varepsilon\tilde{\mathbf{P}} & \mathbf{l} & \mathbf{0} & \mathbf{0} \\ \varepsilon\tilde{\mathbf{P}}' & \mathbf{0} & \mathbf{l}' & \mathbf{0} \\ \varepsilon\tilde{\mathbf{P}}'' & \mathbf{0} & \mathbf{0} & \mathbf{l}'' \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{X}} \\ -\lambda \\ -\lambda' \\ -\lambda'' \end{bmatrix} = \mathbf{0} \qquad (7)$$

The non-zero dimension of the right null-space implies that the 6×6 measurement matrix must have a zero determinant. Expansion of the determinant produces the unique trilinear constraint for 1D views yields

$$\mathbf{l}_i\mathbf{l}'_j\mathbf{l}''_k\mathbf{T}^{ijk} = 0 \qquad (8)$$

$\mathbf{T}^{ijk}$ is the $2 \times 2 \times 2$ homogeneous radial trifocal tensor of the three 1D radial cameras. The expression for the coefficients of the trifocal tensor is

$$\mathbf{T}^{ijk} = \det\left[\tilde{\mathbf{P}}_i^T \tilde{\mathbf{P}}'^T_j \tilde{\mathbf{P}}''^T_k\right] \qquad (9)$$

The radial trifocal tensor is a minimal parameterization of the three $P^2 \rightarrow P^1$ mapping cameras as the d.o.f can be shown to match, $2\times2\times2-1 = 7 = 3\times(2\times3-1)-(3\times3-1)$ (with the LHS being the d.o.f of $\mathbf{T}$ and the RHS being the d.o.f of the three uncalibrated views upto a projectivity) and has no internal constraints.

The radial trifocal tensor can be linearly estimated given seven corresponding triplets (where every triplet gives a linear constraint on the parameters of the radial trifocal tensor using equation (8)) Given more than seven correspondences, we can obtain the linear least squares solution.

It is interesting to verify the relation between the radial trifocal constraint and the radial quadrifocal constraint. When three optical axes intersect, adding a fourth view doesn't yield any additional constraint and the quadrifocal constraint becomes degenerate. Since in $\mathbf{P}^3$ a line and a plane always intersect, we no longer need the precise plane $\mathbf{\Pi}'''$, back-projected from $\mathbf{l}'''$. Instead we could choose any of the planes among the bundle back-projected by the fourth camera. Let us examine the radial quadrifocal constraint, Eq. (4), in this scenario:

$$(\mathbf{l}_i\mathbf{l}'_j\mathbf{l}''_k\mathbf{Q}^{ijk1})l'''_1 + (\mathbf{l}_i\mathbf{l}'_j\mathbf{l}''_k\mathbf{Q}^{ijk2})l'''_2 = 0 \qquad (10)$$

Choosing an arbitrary back-projected plane from the fourth camera corresponds to arbitrary values for l'''_1 and l'''_2. Since Eq. (10) is valid for arbitrary values of l'''_1 and l'''_2, it implies that the coefficients are zero. Further, the above condition is valid for any 3D point $\mathbf{X}$. Comparing this to the trifocal constraint, we see that in this case the quadrifocal tensor must be related to the trifocal tensor as follows

$$\mathbf{Q}^{ijkl} = (\lambda_1\mathbf{T}^{ijk}, \lambda_2\mathbf{T}^{ijk}) \qquad (11)$$

and can only be determined up to one degree of freedom, i.e. $\frac{\lambda_1}{\lambda_2}$.

The trifocal tensor for 1D cameras and its properties were first studied by Quan and Kanade [20] in the context of structure and motion using line correspondences under affine cameras. They showed that by neglecting the position of the lines and considering only their direction, this problem was equivalent to the structure and motion problem for points in one lower dimension. Faugeras et. al. [7] studied the 1D trifocal tensor in the context of planar motion recovery. Our prior work [25] uses the radial trifocal tensor to calibrate cameras with large radial distortion using a polynomial distortion function[3].

[3]This paper only dealt with the trifocal tensor, not the quadrifocal ten-

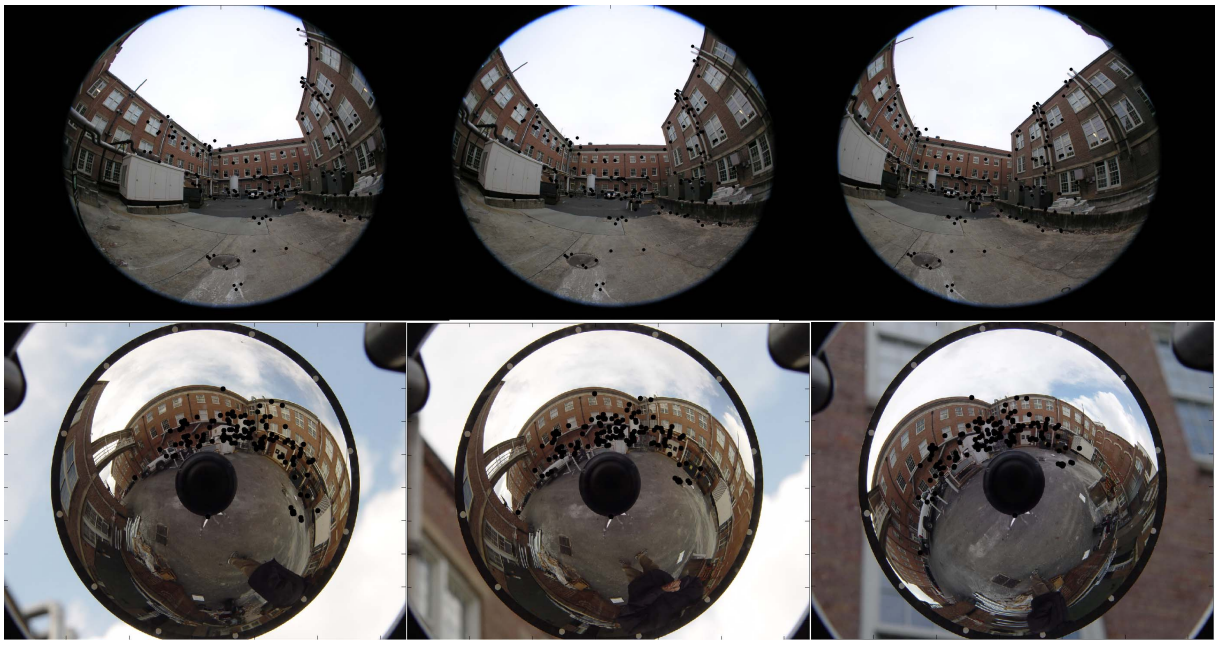

Figure 7. *The triplet of images input to the system with features that were automatically matched overlaid.* ***Top****: Fish-eye images* ***Bottom****: Catadioptric images*

4.1. Reconstruction

4.1.1. Projective Reconstruction

Let us now consider the problem of reconstructing directions from $\mathbf{C}$. Directions correspond to points on Π_∞. Given the radial trifocal tensor, $\mathbf{T}$, we can estimate the three uncalibrated camera matrices, $\tilde{\mathbf{P}}, \tilde{\mathbf{P}}'$ and $\tilde{\mathbf{P}}''$ [25]. These projection matrices can be thought of projecting points on Π_∞ to radial lines in the corresponding views.

Suppose a point on Π_∞, $\mathbf{M}$ (or direction $\mathbf{M}$ in 3D space) projects onto $\mathbf{l}, \mathbf{l}'$ and $\mathbf{l}''$ in the three views. The point can be reconstructed by back-projecting the corresponding radial lines to produce lines on Π_∞. This corresponds to computing the right null-space of the following matrix.

$$\begin{bmatrix} \tilde{\mathbf{P}}^T\mathbf{l} & \tilde{\mathbf{P}}'^T\mathbf{l}' & \tilde{\mathbf{P}}''^T\mathbf{l}'' \end{bmatrix}^T \tag{12}$$

The ray corresponding to features seen in two or more views can be reconstructed.

4.1.2. Metric Reconstruction

Let $\tilde{\mathbf{p}}_1, \tilde{\mathbf{p}}_2$ be the two rows of the projection matrix, $\tilde{\mathbf{P}}_{2\times 3}$. Similarly let $\tilde{\mathbf{p}}'_1, \tilde{\mathbf{p}}'_2$ be the rows of $\tilde{\mathbf{P}}'$ and $\tilde{\mathbf{p}}''_1, \tilde{\mathbf{p}}''_2$ be the rows of $\tilde{\mathbf{P}}''$. Let ω^*_∞ be dual of the absolute conic in the projective frame in which we have reconstructed the points on Π_∞ (directions). It is a 3×3 homogenous symmetric matrix and hence has 5 degrees of freedom (6 upto scale).

sor, and because of the restrictive parametric model this prior work could not model either catadioptric cameras or non-central cameras.

To upgrade the projective reconstruction to metric it is sufficient to estimate ω^*_∞ [13].

We have the assumptions of (i) *zero skew* and (ii) *constant (but possibly unknown) aspect ratio*. It can be shown the assumption of zero skew in the three views gives us the following set of equations linear in the parameters of ω^*_∞ :

$$\begin{aligned} \tilde{\mathbf{p}}_1\omega^*_\infty\tilde{\mathbf{p}}_2^T &= 0 \\ \tilde{\mathbf{p}}'_1\omega^*_\infty\tilde{\mathbf{p}}_2'^T &= 0 \\ \tilde{\mathbf{p}}''_1\omega^*_\infty\tilde{\mathbf{p}}_2''^T &= 0 \end{aligned} \tag{13}$$

Further, the assumption of constant aspect ratio gives us the following equations:

$$\frac{\tilde{\mathbf{p}}_1\omega^*_\infty\tilde{\mathbf{p}}_1^T}{\tilde{\mathbf{p}}_2\omega^*_\infty\tilde{\mathbf{p}}_2^T} = \frac{\tilde{\mathbf{p}}'_1\omega^*_\infty\tilde{\mathbf{p}}_1'^T}{\tilde{\mathbf{p}}'_2\omega^*_\infty\tilde{\mathbf{p}}_2'^T} = \frac{\tilde{\mathbf{p}}''_1\omega^*_\infty\tilde{\mathbf{p}}_1''^T}{\tilde{\mathbf{p}}''_2\omega^*_\infty\tilde{\mathbf{p}}_2''^T} = A \tag{14}$$

If the aspect ratio is known (A is known in Eq. (14)), we have 3 more equations, linear in the parameters of ω^*_∞. If the aspect ratio is unknown, then we have two equations, quadratic in the parameters of ω^*_∞. Using the linear equations in Eq. (13), we can compute the parameters of ω^*_∞ as the intersection of two conics in a plane [4]

4.2. Radial Calibration

Once a metric reconstruction of Π_∞ has been obtained using the 1D radial property of the camera, it can be used to calibrate the remaining unknowns of the projection. In this

[4] The intersection of two conics can be exactly solved analytically

section we present a non-parametric approach to calibrate purely rotating central radially symmetric cameras. Again, this process can be done independently for each image and it is thus possible to calibrate three different cameras using a single trifocal tensor.

For each cameras, the rays corresponding to reconstructed feature points can be represented in a coordinate system relative to the optical axis of the camera, i.e. (θ, ϕ). Because we assume radial symmetry, the ϕ coordinate is irrelevant for us. The goal of our calibration procedure is to obtain a mapping of $\theta(r)$ in function of the radius.

5. Experiments

In our first experiment, a triplet of images obtained using a rotating fish-eye camera was fed as input to the system. The images were acquired using a Nikon 8mm FC-E8 fish-eye converter mounted on a Nikon Coolpix 8400 camera. An online implementation of Lowe's feature matcher was used to obtain triplets of corresponding points in the three images. Note that despite the severe non-perspective distortion, most automatic feature matching techniques work well because the views were obtained using a purely rotating camera.

The image resolution was 1024x768 pixels. Approximately 560 triplets were returned by the feature matcher. Next, RANSAC based on the radial trifocal tensor identified about 220 inliers (the threshold was set to 3 pixels). The input images, with the triplets of corresponding points (those that were identified as inliers after RANSAC) marked are shown in Figure 7. A projective reconstruction was obtained and upgraded to metric based on the assumptions of zero skew and known aspect ratio of unity. For every 3D point that has correspondences across atleast two images, we obtain the angle of the ray, passing through that point and the camera center, and the optical axis of the corresponding view. This gives us a point on the *angle vs. distorted radius* curve. Figure 8 shows the plots for the three views. We see that the angle of a ray with the optical axis is related to the distorted radius almost linearly. This is expected as a fish-eye camera roughly follows the equidistant model.

Note that at no point during the whole procedure did we make any assumptions about the type/amount of radial distortion. Further, an automatic feature matcher has been able to give us features that span the whole range of distorted radii. Finally, note that no additional constraint (smoothness etc) was enforced across the three views. Finally in Figure 9 we show a cubemap of the undistorted left view. Note that straight lines in the world are indeed mapped to straight lines in the image. The unwarping was carried out by computing the distorted radius for a given undistorted radius using a simple line interpolation on the plot in Figure 8.

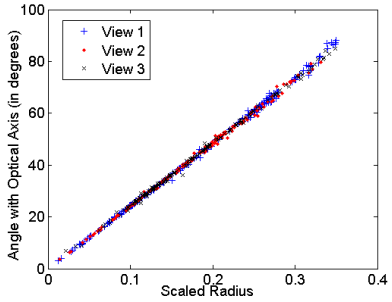

Figure 8. *Plot for fish-eye camera.*

More complex models could also be used .

In our second experiment, three images obtained from a purely rotating single viewpoint catadioptric camera were used. The image resolution was 1280x960 pixels. Lowe's feature matcher [17] produced approximately 220 matching triplets across the three views. An in the previous experiment, RANSAC based on the radial trifocal tensor produced around 130 inlier triplets. Figure 7 shows the input images with the inlier triplets marked. We compute a projective reconstruction and upgrade it to metric based on assumptions of zero skew and unit aspect ratio. Note that since our method handles all types of radial distortion uniformly, the complete procedure in this experiment in exactly the same as in the previous experiment. Figure 10 shows the plots of the *Angle with Optical Axis vs. Distorted Radius* for each of the views. Finally, Figure 11 shows a cubemap of the undistorted left view. One can refine the estimates produced by our method using techniques like bundle adjustment. In the first experiment, it reduced the RMS reprojection error from 1.13 pixels to 0.43 pixels.

Figure 9. *Cubemap of undistorted left image*

6. Conclusion

In this paper we have introduced the 1D radial camera which maps 3D points on radial lines. This allows us to derive multilinear constraint between three and four views recorded with central or non-central omnidirectional cameras. Given 15 or more correspondences across four views

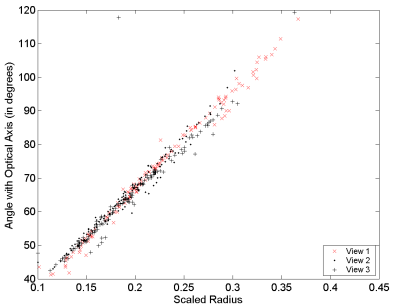

Figure 10. *Plot for catadioptric camera.*

Figure 11. *Cubemap of undistorted left image*

taken with a moving camera (or multiple cameras), the corresponding 1D radial cameras and a metric reconstruction of 3D points can be computed. We have then demonstrated how the reconstruction can be used as calibration object to estimate non-parametric camera model for different cameras, including a non-central cameras. Although the approach has been demonstrated to work well on synthetic data, the required number of point correspondences makes it hard to develop a robust automatic approach for real images.

For a purely rotating camera a simpler constraint is obtained requiring 7 point correspondences across 3 views. As with perspective cameras [11, 5], pure rotation turns out to be particularly well suited for self-calibration of central omnidirectional cameras. In particular, we present an automatic approach that recovers the accurate non-parametric distortion curve relating image radius to angle of incoming rays. We use our approach to automatically calibrate both fish-eye lenses and cata-dioptric cameras from real images.

References

[1] D. Aliaga. Accurate catadioptric calibration for real-time pose estimation of room-size environments. In *ICCV*, 2001.

[2] S. Baker and S. Nayar. A theory of single-viewpoint catadioptric image formation. *IJCV*, 35(2):1 – 22, 1999.

[3] H. Bakstein and T. Pajdla. Panoramic mosaicing with a field of view lens. In *Proc. IEEE Workshop on Omnidirectional Vision*, pages 60–67, 2002.

[4] R. Benosman and S. Kang. *Panoramic Vision: Sensors, Theory and Applications*. 2001.

[5] L. d. Agapito, R. Hartley, and E. Hayman. Linear selfcalibration of a rotating and zooming camera. In *CVPR*, 1999.

[6] F. Devernay and O. Faugeras. Straight lines have to be straight. *MVA*, 13(1):14–24, 2001.

[7] O. Faugeras, L. Quan, and P. Sturm. Self-calibration of a 1d projective camera and its application to the self-calibration of a 2d projective camera. *PAMI*, 22(10):1179–1185, 2000.

[8] C. Geyer and K. Daniilidis. Structure and motion from uncalibrated catadioptric views. In *CVPR*, 2001.

[9] C. Geyer and K. Daniilidis. Paracatadioptric camera calibration. *PAMI*, 24(5):687–695, May 2002.

[10] M. Grossberg and S. Nayar. A general imaging model and a method for finding its parameters. In *ICCV*, 2001.

[11] R. Hartley. Self-calibration from multiple views with a rotating camera. In *ECCV*, pages 471–478, 1994.

[12] R. Hartley and F. Schaffalitzky. Reconstruction from projections using grassmann tensors. In *ECCV*, 2004.

[13] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. 2000.

[14] R. A. Hicks and R. Bajcsy. Catadioptric sensors that approximate wide-angle perspective projections. In *CVPR*, 2000.

[15] S. Kang. Catadioptric self-calibration. In *CVPR*, 2000.

[16] J. Kannala and S. Brandt. A generic camera calibration method for fish-eye lenses. In *ICPR*, 2004.

[17] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.

[18] B. Micusik and T. Pajdla. Estimation of omnidirectional camera model from epipolar geometry. In *CVPR*, 2003.

[19] M. Pollefeys, R. Koch, and L. V. Gool. Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters. In *IJCV*, 1999.

[20] L. Quan and T. Kanade. Affine structure from line correspondences with uncalibrated affine cameras. *PAMI*, 1997.

[21] S. Shah and J. Aggarwal. Intrinsic parameter calibration procedure for a (high-distortion) fish-eye lens camera with distortion model and accuracy estimation. *Pattern Recognition*, 29(11):1175–1788, 1996.

[22] P. Sturm and S. Ramalingam. A generic concept for camera calibration. In *ECCV*, volume 2, pages 1–13, 2004.

[23] R. Swaminathan, M. Grossberg, and S. Nayar. Caustics of catadioptric camera. In *ICCV*, volume 2, pages 2–9, 2001.

[24] R. Swaminathan and S. Nayar. Nonmetric calibration of wide-angle lenses and polycameras. *PAMI*, 2000.

[25] S. Thirthala and M. Pollefeys. The radial trifocal tensor: A tool for calibrating the radial distortion of wide-angle cameras. In *CVPR*, 2005.

[26] B. Triggs. Matching constraints and the joint image. In *ICCV*, pages 338–343, 1995.

[27] B. Triggs. Autocalibration and the absolute quadric. In *CVPR*, pages 609–614, 1997.

[28] R. Willson and S. Shafer. What is the center of the image? In *CVPR*, pages 670 – 671, June 1993.

[29] Y. Xiong and K. Turkowski. Creating image-based vr using a self-calibrating fisheye lens. In *CVPR*, pages 237–243, 1997.

[30] X. Ying and Z. Hu. Catadioptric camera calibration using geometric invariants. In *ICCV*, 2003.