On the Global Self-calibration of Central Cameras Using Two Infinitesimal Rotations

Ferran Espuny

Université Paris Descartes, Laboratoire MAP5 (CNRS UMR 8145)
Ferran.Espuny@parisdescartes.fr

Abstract. The calibration of a generic central camera can be described non-parametrically by a map assigning to each image pixel a 3D projection ray. We address the determination of this map and the motion of a camera that performs two infinitesimal rotations about linearly independent axes. A complex closed-form solution exists, which in practice allows to visually identify the geometry of a range of sensors, but it only works at the center of the image domain and not accurately.

We present a new two-step method to solve the stated self-calibration problem that overcomes these drawbacks. Firstly, the Gram matrix of the camera rotation velocities is estimated jointly with the Lie bracket of the two rotational flows computed from the data images. Secondly, the knowledge that such Lie bracket is also a rotational flow is exploited to provide a solution for the calibration map which is defined on the whole image domain. Both steps are essentially linear, being robust to the noise inherent to the computation of optical flow from images.

The accuracy of the proposed method is quantitatively demonstrated for different noise levels, rotation pairs, and imaging geometries. Several applications are exemplified, and possible extensions and improvements are also considered.

1 Introduction

There exists a wide variety of central cameras, not limited to systems with multiple lenses (dioptric), but also comprising combinations of lenses and mirrors (catadioptric) [1]. Rather than using a specific parametric model for each central system, we consider the non-parametric generic camera model, which calibration map associates to each image pixel a projection ray [2,3,4].

Concisely, we address in this paper the self-calibration of this model (estimation of both camera motion and calibration map) using only images acquired with two infinitesimal rotations of the camera. The considered problem is closely related to that of recovering a mirror's shape from rotational specular flows [5]. However, rotational specular flows offer additional cues for self-calibration [6].

In the case of central cameras, the generic calibration problem (finding a projection direction for each pixel) is equivalent to non-parametric distortion correction for those cameras with a field of view smaller than 180 degrees; for this equivalence, it suffices to place an auxiliary plane in front of the camera. Other

K.M. Lee et al. (Eds.): ACCV 2012, Part IV, LNCS 7727, pp. 243–256, 2013.

Fig. 1. Left: example image acquired by a central camera with high radial distortion, and superimposed the optical flows computed using two small camera rotations. Centre: undistortion obtained with the state-of-the-art method in [18]. Right: global undistortion obtained with our non-parametric linear method; its average global error is 2 pixels, from which only 0.5 pixels correspond to non-perspective distortion.

applications of camera calibration include motion estimation and mosaicing/3D reconstruction from images. Examples are shown through the paper.

Several non-parametric solutions exist using a calibration pattern for (generic) camera calibration [2,4] and distortion correction [7,8]. The self-calibration of non-parametric distortion models, limited to be radial, has been addressed in [7,9,10,11]. In contrast, we do not assume any structure on the scene geometry, and only require the generic calibration map to be smooth.

The metric self-calibration (i.e. up to rotation) of a generic central camera has been solved, at least at a proof-of-concept level, for particular (non-infinitesimal) motions [12,13,14]. However, these methods estimate *motion flows* from image matches and thus the camera relative rotations can not be general. In contrast, we use infinitesimal camera rotations with axes being allowed to be general.

Another generic self-calibration method exists that requires a large image dataset to (inaccurately) estimate a discrete generic calibration map, assumed to be correlated with dissimilarity measures in the input images [15]. In contrast, we only require two rotational image sequences to solve a problem equivalent to theirs in the case of smooth generic central cameras.

In [16], the optical flows produced by three infinitesimal camera rotations are used for the self-calibration, up to projective transformation, of a smooth generic central camera. In [17], two rotations are used for the metric self-calibration of said camera model, the motion estimation depending on second order derivatives of the data flows, and the calibration map estimation depending rationally on the flow coordinates. A re-formulation of [17] using the Lie bracket of the flows is given in [18]; results corresponding to "real" data flows are shown, which are neither accurate nor defined near the image borders.

The computation of rotational flows compatible with the existence of a common camera is outlined in [18] as a potential improvement of [17,18]; however, it is cast as "highly non-linear optimization problems", which remain unsolved. It is also shown in [18] that three different rotations may be used to achieve a global result, more regular, but still inaccurate, even if using exact motions.

We present a new step-wise linear method for the global self-calibration of a smooth generic central camera from two rotational flows (see Fig. 1). We show that the camera motion can be accurately determined by its joint estimation with the Lie bracket of the two data flows, which in fact corresponds to a third unobserved rotational flow. This first step represents a workaround solution to the computation of compatible rotational flows pursued in [18]; moreover, it requires the use of derivatives of the data flows only up to first order.

We also show that the Lie bracket flow obtained during the motion estimation step allows an accurate global estimation of the calibration map. This second computation overcomes the division by the data flow coordinates associated to the closed-form formulae in [17,18]: the calibration map can be linearly determined thanks to the previous estimation of the Lie bracket flow.

Next, we introduce the necessary notation and background, concisely stating the problem; in Section 3, we overview and analize the existing closed-form solution. We present our two-step method in Section 4. An experimental evaluation and analyisis of our proposal is performed in Section 5 using simulated image sequences, and an example with real images is also shown before the Conclusion.

Notation. In order to allow an easy comparison with [18], we adopt their notation conventions. The symbol ∂ will be used for differentiation, a cross symbol $\times$ will denote the cross product operator, and Id_2 the identity matrix of size 2.

2 A Generic Self-calibration Problem

2.1 Preliminaries

Following [2,3,4] we consider a *generic camera* to be a set of image points in (possibly non-parametric) correspondence with a set of 3D projection rays. We say that a generic camera is *central* if all its projection rays intersect in a single point, called the *camera centre* [2,3,4]. In this case, we can use the unit sphere S^2 to describe the possible projection rays. Accordingly, we define the *calibration map* of a generic central camera as a map from $\mathcal{U}$, an open connected subset of $\mathbb{R}^2$ (image pixels), on the unit sphere S^2 (oriented projection directions):

$$\begin{array}{rcl} f : \mathcal{U} \subset \mathbb{R}^2 \rightarrow & & S^2 \\ (u, v) & \mapsto & f(u, v) \; . \end{array} \tag{1}$$

This map sends a planar image to its undistorted version on the sphere. In order to use optical flow, we assume that the calibration map f is smooth [16,17,18].

In the following, we will use f for theoretical demonstration purposes, with no further constraint on it that having norm one. In contrast, we will only show results corresponding to calibration maps with $f_3 > 0$, meaning in practice that the camera angular field of view is smaller than 180 degrees (for omnidirectional cameras, the visualization of a calibration map is not straightforward). For this purpose, we introduce the *undistortion map*

$$g = (g_1,\ g_2)^T := (f_1/f_3,\ f_2/f_3)^T \; . \tag{2}$$

In Fig. 1 we show an example of image with high radial distortion, together with its undistorted version using the map g. By the smoothness assumption on the calibration map f, the undistortion map g is also smooth, which we will impose using b-splines (see the Appendix).

We consider the *optical flow* in a sequence of images to be the velocity field in the image domain tangent to the image transformation that takes one image into the next one (not the transformation itself).

2.2 Problem Statement

Assume that we know two point-wise linearly independent optical flows V_1, V_2 observed on an open subset $\mathcal{U}$ of the image, corresponding to rotations of a generic central camera about two linearly independent axes passing through the camera centre. The self-calibration problem consists in determining the camera rotational velocities ω_1, ω_2 and the calibration map f that are compatible with these flows, i.e. satisfying the following equations [3,18]:

$$\mathrm{D}f(u,v) \cdot V_i(u,v) = -\omega_i \times f(u,v) \; . \tag{3}$$

It is proven in [17] that this problem can be solved up to an orthogonal transformation. Accordingly, we assume as given an orthonormal basis $\{u_1, u_2\}$ so that

$$(\omega_1, \omega_2) = (u_1, u_2) \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} \; , \tag{4}$$

for some unknown scalars a, b, c satisfying $a, c > 0$. Observe that this introduces an asymmetry in the problem, since the direction of one of the axis of rotation is known, whereas the other axis is only constrained to lie on a known two-dimensional semi-space. A symmetric but less general approach, not followed here, would be to consider as given the two directions of rotation.

3 Existing Closed-Form Solution and Analysis

The generic self-calibration of a central camera from two rotational flows is solved in closed-form in [17,18], using two auxiliary functions given by:

$$\Delta_1 = -\mathrm{tr}(\mathrm{D}V_2) + \frac{1}{\det V} \mathrm{D} \det V \cdot V_2 \; , \tag{5}$$

$$\Delta_2 = \quad \mathrm{tr}(\mathrm{D}V_1) - \frac{1}{\det V} \mathrm{D} \det V \cdot V_1 \; . \tag{6}$$

Concisely, the Gram matrix

$$G_\omega := (\omega_1, \omega_2)^T \cdot (\omega_1, \omega_2) \tag{7}$$

is determined by averaging the following pixel-wise estimators, taking into account the expected positive definiteness of G_ω:

$$G_\omega = \begin{pmatrix} \Delta_2 \\ -\Delta_1 \end{pmatrix} \cdot (-\Delta_2, \Delta_1) + \begin{pmatrix} \mathrm{D}\Delta_2 \\ -\mathrm{D}\Delta_1 \end{pmatrix} \cdot V \; . \tag{8}$$

Then, the camera motion is extracted from this Gram matrix using (4) and (7), and finally the calibration map f is computed as the norm-one map such that

$$f \propto \left(\omega_1, \omega_2, \omega_1 \times \omega_2\right) \cdot \left(\Delta_1, \Delta_2, 1\right)^T \ . \tag{9}$$

Remark 1 (Analysis). By (9), the relation between the undistortion map g, defined by (2), and the functions Δ_1, Δ_2, defined by (5–6), is rational except when ω_1 and ω_2 are orthogonal to the Z axis. For instance, for $\omega_1 = (0, 1, 0)^T$ and $\omega_2 = (0, 0, 1)^T$, we have that $g_1 = 1/\Delta_2$, $g_2 = \Delta_1/\Delta_2$; therefore, any error in the estimation of Δ_2 affects severely the estimation of the undistortion map (g_1, g_2). In addition, the formulae (5–6) require the first order derivatives of the data flows, which are later differentiated in (8) to estimate the camera motion. As a result, the method in [17,18] has trouble to determine the solution close to the image borders in presence of noise or real flows.

It is shown in [18] that the functions (Δ_1, Δ_2) defined by (5–6) are in fact (with reversed sign) the coordinates in the point-wise vector basis V_1, V_2 of the Lie bracket vector field:

$$[V_1, V_2] := \mathrm{D}V_2 \cdot V_1 - \mathrm{D}V_1 \cdot V_2 \ . \tag{10}$$

The geometric interpretation of this Lie bracket is also given: it is the optical flow of a (non-performed) rotation with angular velocity $\omega_1 \times \omega_2$. As a consequence, Eq. (8) can be written as:

$$G_\omega = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \cdot (V_1, V_2)^{-1} \cdot \left([V_1, [V_1, V_2]] \ \ [V_2, [V_1, V_2]]\right) \ . \tag{11}$$

In the next section, we exploit these theoretical formulae from [18] to propose a new global self-calibration method.

4 A Two-Step (Linear) Method

Assume that we are given two optical flows V_1, V_2 as described in Section 2.2. These flows can be computed linearly from two initial image sequences by minimizing the Linearized Brightness Constancy Constraint [20], which we do in practice by following [18] with a b-spline model. We denote the coefficients of the Gram matrix G_ω in (7) as $G_{i,j} := \omega_i^T \cdot \omega_j$.

4.1 Estimation of the Camera Angular Velocities and the Lie Bracket Flow

Due to noise in V_1, V_2, a direct computation using (10) of the Lie bracket flow

$$V_3 = [V_1, V_2] \tag{12}$$

is not likely to be compatible with the desired existence of a constant positive definite 2×2 matrix G_ω satisfying (11). Since the optical flows V_1, V_2 are assumed to be known, such constraint can be written as a differential linear combination of V_3 and $G_{i,j}$:

$$[V_1, V_3] + G_{1,1}V_2 - G_{1,2}V_1 = 0\,, \tag{13}$$
$$[V_2, V_3] + G_{1,2}V_2 - G_{2,2}V_1 = 0\,. \tag{14}$$

In summary, denoting by $V_i^T \cdot \nabla$ the operator $V_{i,1}\partial_u + V_{i,2}\partial_v$, the flow V_3 and the coefficients $G_{i,j}$ of the Gram matrix satisfy the following 3 vector equations (i.e. 6 scalar equations) per pixel:

$$\begin{pmatrix} \mathrm{Id}_2 & 0 & 0 & 0 \\ \mathrm{Id}_2\left(V_1^T \cdot \nabla\right) - \mathrm{D}V_1 & V_2 & -V_1 & 0 \\ \mathrm{Id}_2\left(V_2^T \cdot \nabla\right) - \mathrm{D}V_2 & 0 & V_2 & -V_1 \end{pmatrix} \cdot \begin{pmatrix} V_3 \\ G_{1,1} \\ G_{1,2} \\ G_{2,2} \end{pmatrix} = \begin{pmatrix} [V_1, V_2] \\ 0 \\ 0 \end{pmatrix}\,. \tag{15}$$

In practice, we model the smooth optical flow V_3 using b-splines, as explained in the Appendix. The resulting sparse linear system has $2N + 3$ unknowns, being N the size of the b-spline coefficient vector. If we do not impose the positive definiteness on G_ω, it can be solved either using least squares or a more robust procedure, as detailed in Appendix. Otherwise, the corresponding constrained problems can be solved using Second Order Cone Programming (SOCP) [19]. Given that the order of the Lie bracket coefficients can be quite dissimilar from that of the Gram matrix coefficients, we use an initial estimation given by the L_2 optimization to normalize the equations and then re-estimate the parameters.

Remark 2 (Potential Use for Parametric Camera Rotational Flows). Equation (15) does only involve the optical flows and the camera motion: it is independent of the model used to describe a central camera. Therefore, it can be used in a (possibly uncalibrated) parametric context as a constraint on rotational flows and/or on the motion of a rotating camera using those flows for computation and/or evaluation purposes.

Remark 3 (Computing Optical Flows Compatible with the Rotation of a Common Camera). A further improvement of the previous motion estimation method could be achieved by the joint estimation, directly from the rotational image sequences, of the two rotational flows, their Lie bracket and the Gram matrix of the camera motions. The difficulty in such problem lies in the need for using a robust penalization for the non-linear terms arising from (13),(14) when considered as a function of V_1, V_2, V_3, $G_{i,j}$, and shall be a topic of future research (experiments using L_2 penalization were not successful).

4.2 Linear Estimation of the Global Undistortion Map

After the previous step, we may assume as known both the rotational flows V_1, V_2, their Lie bracket V_3, and the corresponding camera angular velocities ω_1,

ω_2, $\omega_1 \times \omega_2$, which can be determined with the Cholesky decomposition of the Gram matrix G_ω using (7).

Observe that now the joint matrix of the three available rotational flows, (V_1, V_2, V_3), has point-wise rank two (except at most in two isolated image pixels). Therefore, we can compute the undistortion map g linearly by adapting the 3-flow methods in [16,18]. Concisely, using that, by (4), we have

$$(\omega_1, \omega_2, \omega_1 \times \omega_2)^{-1} = \frac{1}{ac} \begin{pmatrix} cu_1^T - bu_2^T \\ au_2^T \\ {u_1 \times u_2}^T \end{pmatrix} , \tag{16}$$

the undistortion map (g_1, g_2) satisfies the following linear constraint:

$$\big(V_1, V_2, V_3\big) \cdot \begin{pmatrix} cu_1^T - bu_2^T \\ au_2^T \\ u_1 \times u_2^T \end{pmatrix} \cdot \begin{pmatrix} g_1 \\ g_2 \\ 1 \end{pmatrix} \;=\; 0 \,. \tag{17}$$

In practice, we model the undistortion map g using b-splines, as explained in the Appendix, and we solve the resulting sparse linear system with either L_2 or robust $L_{1\varepsilon}$ penalization, as detailed in that section. Observe that the result (17) still holds true if we take a general calibration map f instead of $(g_1, g_2, 1)^T$, and therefore it may be used for omnidirectional cameras.

4.3 Summary

We conclude this section by summarizing the proposed self-calibration method, leaving clear the used parameters and its computational cost. In contrast with the existing closed-form method [17,18], not only Step 1 but all the computations required by the algorithm are step-wise linear, excepting a Cholesky decomposition of the Gram matrix G_ω in Step 3, and a positive constraint on G_ω and a quadratic formula for $[V_1, V_2]$ in Step 2. Moreover, with respect to that early method, the use of flow derivatives has been reduced to first order, and these are only needed for computing the Lie bracket flow $V_3 = [V_1, V_2]$ in Step 2.

Algorithm 1. Self-calibration from Two Infinitesimal Rotations

Input. Two sequences of rotational images; two orthonormal 3-vectors u_1, u_2
1. Compute the generic rotational flows V_1, V_2 from the images
2. Compute $V_3 = [V_1, V_2]$ and G_ω by solving (15)
3. Extract the rotation angular velocites ω_i from G_ω and the u_i using (4)
4. Compute g by solving (17)

return ω_1, ω_2, g

Remark 4 (Number of Parameters). The proposed algorithm only requires as parameters two orthonormal vectors to avoid the orthogonal ambiguity in the self-calibration problem (Section 2.2). In practice, we use 2D b-splines in several estimations: the optical flows V_1, V_2 in Step 1, their Lie bracket V_3 in Step 2,

and the undistortion map g in Step 4. Since we assume equi-distributed knots (Appendix), we need three parameters for each b-spline: the number of knots in each direction, n_1, n_2, and a smoothing factor λ. We only used the latter when computing the optical flows in Step 1, where an extra parameter σ is also needed for pre-smoothing the images.

Remark 5 (Computational Cost). The algorithm requires the resolution of linear systems of different sizes: $n \times 2N$ in Step 1, $6n \times (2N+3)$ in Step 2, and $2n \times 2N$ in Step 4, being n the number of image pixels and N the length of each b-spline coefficient vector. We solve both the L_2 minimization (18) and the iteratively re-weighted $L_{1\varepsilon}$ minimization (19) through the computation of the corresponding normal equations and a Cholesky-based resolution of those equations. When solving (15), we only impose the positive definiteness on the Gram matrix G_ω and solve the constrained least squares problem if the unconstrained linear methods fail to find a positive definite solution; in practice, this will likely happen with either high noise or with optical flows (close to) being linearly dependent.

5 Experimental Results

5.1 Error Measures

For simulated image sequences, we will measure the errors in optical flow (using angular and relative norm errors) and the camera angular velocities (using relative metric errors) as in [17,18]. In addition, we introduce two measures for the evaluation of an estimated undistortion map g when its groundtruth $\hat{g}$ is known. First, the absolute global error GE (in pixels), defined pixel-wise by $GE = \text{dist}(g, \hat{g})$. This serves as an overall error, since the calibration results are coupled with the motion estimation errors, as it follows from formula (17).

We use as second error measure of an estimated g the error after correcting it with the "best" homography H approaching g to $\hat{g}$, H being computed with the DLT algorithm. We refer to this as the absolute non-perspective error, NPE (in pixels), defined at each pixel as $NPE = \text{dist}(H \cdot g, \hat{g})$. The motivation for this choice is that the correctly undistorted images should be perspective-like images, for which the infinitesimal camera rotations induce image homographies.

5.2 An Example with High Radial Distortion

We generated sequences of 500×500 images corresponding to rotations about the Y and Z axes, with angular velocities of norm equal to 0.003, of a camera with high radial distortion (Fig. 1 left contains an example image). To avoid "maquillaging" the solution, we fixed beforehand the b-spline parameters $n_1 = n_2 = 5$ and $\lambda = 0$ in (19) for the computation of the Lie bracket V_3 or the undistortion map g.

We computed the optical flows V_1, V_2 corresponding to only the first two images of each sequence by minimizing the $L_{1\varepsilon}$ penalization of the Linearized Brightness Constancy Constraint [20] with $\sigma = 3.0$ for image pre-smoothing, and

respective thin-plate regularity weights $\lambda_{V_1} = 10^6$, $\lambda_{V_2} = 10^{10}$ (we refer to the Appendix for this parameter). The groundtruth flows are depicted in Fig. 1; it can be observed that the flow V_1 (horizontal green arrows) has a very changing scale according to the different resolutions in the image. Due to this fact, in general the selection of a constant smoothing parameter σ is not optimal, and we expect bigger errors, specially in the derivatives, for high scale changing flows. As we show later, image smoothing can be improved after calibration information is available. The results are summarized in Table 1 (left). We include the errors for the flow derivatives because they are used for the next motion estimation step, and they turn out to be much higher than the flow errors for V_1.

Table 1. Left: Average and standard deviation of the angular error AE (in degrees) and relative norm error RNE (%) of the data optical flows and derivatives. Right: relative metric errors (%) in the estimated camera angular velocities, and errors in the estimated Lie bracket flow, both being used in the second self-calibration step.

Flow	μ_{AE}	σ_{AE}	μ_{RNE}	σ_{RNE}
V_1	0.18	0.19	0.90	0.59
$\partial_u V_1$	3.19	3.49	3.29	3.47
$\partial_v V_1$	2.62	2.57	1.50	1.50
V_2	0.13	0.21	0.25	0.31
$\partial_u V_2$	0.08	0.04	0.04	0.03
$\partial_v V_2$	0.06	0.03	0.08	0.06

	Motion Errors			**Lie Bracket Errors**			
Method	$\|\omega_1\|$	$\|\omega_2\|$	$\widehat{\omega_1,\omega_2}$	μ_{AE}	σ_{AE}	μ_{RNE}	σ_{RNE}
Esp07 [17]	4.13	0.14	1.86	1.28	1.91	2.11	2.88
E_{L_2}	2.30	0.39	1.82	0.53	0.54	1.78	2.49
$E_{L_{1\varepsilon}}$	**0.39**	**0.02**	**0.54**	**0.21**	**0.20**	**0.83**	**0.46**

Given the two rotational flows, we computed their Lie bracket and the two camera angular velocities using the three available methods. First, the closed-form formulae and the averaging process in [17,18]; second (resp. third), their joint optimization as described in Section 4.1 with a L_2 (resp. $L_{1\varepsilon}$) penalization for the resulting linear system (15). We see from the results, summarized in Table 1 (right), that the $L_{1\varepsilon}$ method overperforms the other ones: it estimates the motion with far below a 1% of relative error.

Fig. 2. Left: groundtruth sensor (red lines) and the estimated undistortion map (dashed blue lines). Centre: original image smoothed using the metric induced by this map. Right: mosaic performed with the Z rotational sequence of calibrated images.

We finally computed the calibration map using the $L_{1\varepsilon}$ penalization, and obtained results as observed in Fig. 1: the global error had mean and standard deviation values $\mu_{GE} = 2.548, \sigma_{GE} = 0.758$ (in pixels), whereas the non-perspective correction error was $\mu_{NPE} = 0.568, \sigma_{NPE} = 0.595$ (in pixels), the error measures being computed as explained at the beginning of this section. It can be observed (better zooming in) at Fig. 2 left that, as typical in methods for the correction of geometric distortion, the calibration errors are mostly concentrated at the image corners. In comparison with the state of the art in [18], ours is a more accurate and globally defined solution.

The achieved accuracy allows us to perform typical operations with calibrated cameras. A first example of application consists in smoothing the original images taking into account the sensor geometry, i.e. non-uniformly (Fig. 2, centre). This is done by solving the heat-diffusion equation [21], which we perform via an iterative evaluation of the Laplace-Beltrami operator and an update of the smoothed image, as for the parametric spherical camera model in [22].

A second example application, using the obtained undistorted images and motion estimation, consists in creating a mosaic by reversing the effect of the estimated rotations to place all such images on a common frame. In Fig. 2 right we show a 560×560 mosaic generated in this way using a sequence of 50 images corresponding to a camera rotation about the Z axis, the first two images being the ones used for camera calibration.

5.3 Further Evaluation

It seems natural to ask what would have been the ouput in the previous experiment if the noise in the flows had been different from the obtained with the given images. Moreover, we may wonder whether the two measured flow errors, namely the angular and relative norm error, affect the results in equal measure or not; we study these two factors as possible main error sources, although, as already pointed out, our results depend not only on the optical flow components but also on their derivatives.

In Fig. 3 we show the average result errors after performing 50 simulations of noise in the two data flows for each combination of angular and relative norm errors between 0.2 and 1. We fitted b-splines to the noisy data, again without tuning the involved parameters, being the actual average noise in the fitted b-spline flows possibly higher than reported. This simulation must be considered carefully, taking into account that two rotational flows in general do not have the same angular and norm noise levels; moreover, it is mostly tentative, since other factors affect the calibration. In all the tested cases, we outperformed [18] and obtained the worst results for bigger norm errors in the data flows, being the motion estimation also affected, in smaller measure, by the angular errors.

We performed a last evaluation experiment by fixing the Y axis and varying the second rotation axis (Fig. 4). We obtained the best results for the axes orthogonal to the Y axis. The results are again sensor-dependent (for the used radial distortion sensor, the Z axis optical flow is quite easy to model), and could be improved by using the smoothness parameters that we set to zero.

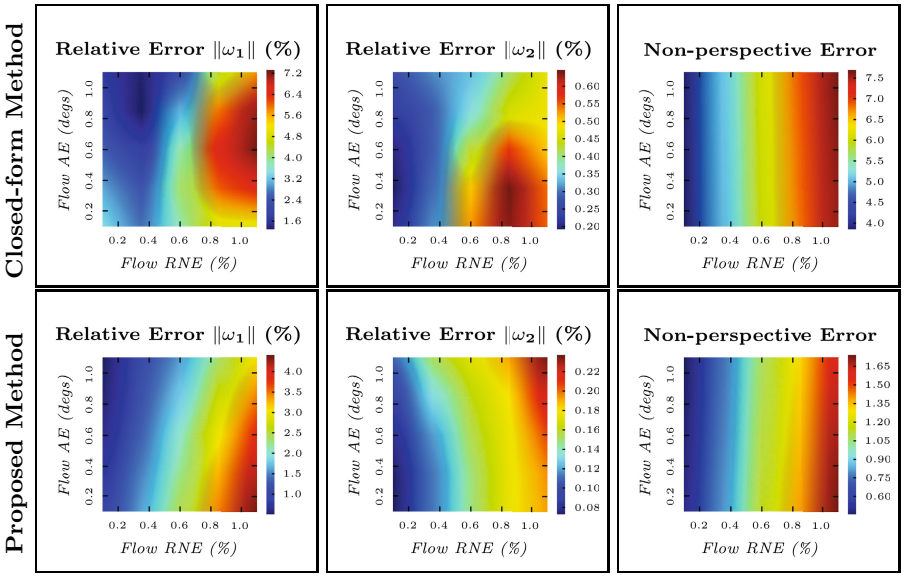

Fig. 3. Errors corresponding to the method in [18] (top row) and ours (bottom row) for the Y and Z rotations with different levels of relative norm and angular error in the two data optical flows. We omitted: the errors in $\widehat{\omega_1, \omega_2}$, which were similar to those for $\|\omega_2\|$, and the global errors, which were mostly affected by the bigger errors in $\|\omega_1\|$.

5.4 A Highly Non-linear Sensor

In Fig. 5, we consider a sinusoidal sensor inspired by [16,18], which is maybe not very realistic, but it shows the power of our non-parametric method to self-calibrate highly non-linear sensors. Taking only as input two sequences of 10 images (Y and Z axis rotations), our method estimates the camera motion with relative metric errors below 0.5%. The errors, explained at the beginning of this section, without imposing regularity are quite good: $\mu GE = 1.24$, $\mu NPE = 0.82$.

5.5 Real Images

We finally consider the fish-eye image sequences from Section 8.2 in [18].[1]The image sequences are particularly cumbersome for optical flow computation, due to the lack of smoothness or texture in certain image regions; our results, shown in Fig. 6, are clearly worst in those regions. The mean error between the corners of the undistorted checkerboard pattern and the projection of an ideal pattern under a best least-squares fitting homography is the 17.1% of the average length of an undistorted square side (standard deviation equal to 9.8). A further bundle adjustment process is currently under study. The interested reader may compare our results with those in Fig. 10 in [18] to assess the improvement in the area of image that we calibrate (improvement already outlined in Fig. 1).

[1] Images available at
`http://atlas.mat.ub.es/personals/fespuny/Research.html#CGRF`

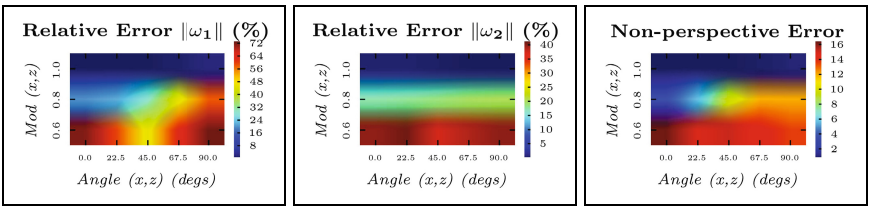

Fig. 4. Average errors obtained with AE=0.15 degs, RNE=0.5% in two flows corresponding to the Y axis and different linearly independent second axes; we use polar representation in the (x,z) coordinates of the second axis, e.g. angle 0 and norm 1 means the X axis, and angle 90 and norm 0.6 represents the axis of direction $(0, 0.8, 0.6)$.

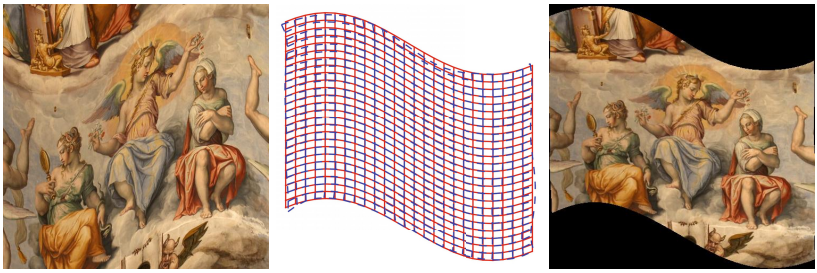

Fig. 5. Left: example image with sinusoidal distortion. Centre: estimated undistortion map (dashed blue lines). Right: undistorted image using b-spline regularity ($\lambda = 10^3$).

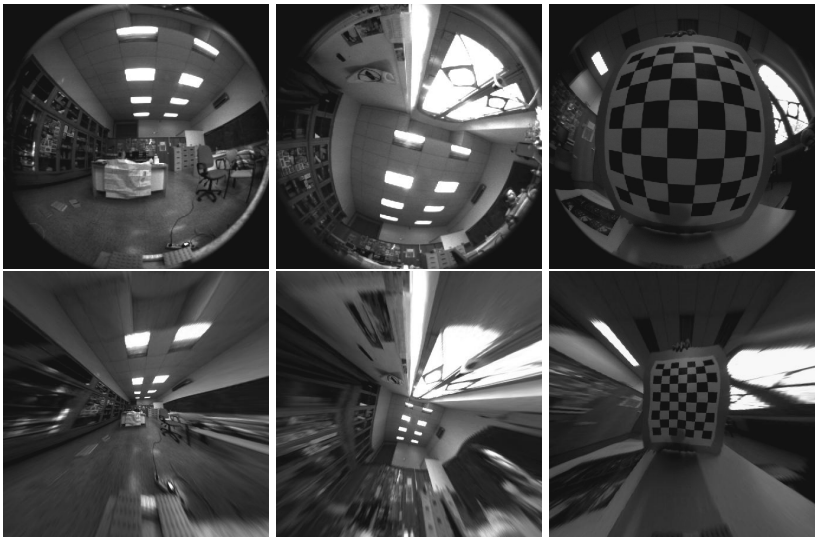

Fig. 6. Top: original images from [18]. Bottom: global undistortion results.

6 Conclusion

We have presented a new method for the global self-calibration of central cameras using the optical flows produced by two infinitesimal camera rotations. Despite a positive definiteness constraint on the Gram matrix of the rotation angular velocities, all the involved steps are linear, being the method quite robust and accurate, specially for orthogonal rotation axes. As discussed in the paper, the results can be applied both in parametric and non-parametric settings.

In fact, we have shown that any two optical flows covering an image region give enough information for self-calibrating that region. Therefore, when having more than two rotations, the resulting optical flows are no longer required to be dense for self-calibrating the camera, as far as they locally overlap pair-wise on the whole image. In conclusion, the simplicity of our proposal may conduct to a highly dynamical and accurate non-parametric method for the self-calibration of central cameras with multiple infinitesimal rotations.

Acknowledgement. The author is grateful to José I. Burgos Gil, Lionel Moisan, and Pascal Monasse for valuable discussions. This work was partly supported by MAP5 (CNRS UMR 8145) and the projects MTM2009-14163-C02-01 (Spain), Callisto ANR-09-CORD-003 (France), and IMAGINE-ENPC (France).

Appendix. Using B-Splines for Smooth 2D Functions

We are interested in the estimation of several two-dimensional smooth functions $F = F(u, v) = (F_1, F_2)^T$ defined on the image, e.g. the undistortion map in (2). We will model each component F_i as a tensor b-spline surface with a grid of $n_1 \times n_2$ equi-distributed knots (see e.g. [23]): $F_i(u, v) = w(u, v)^T \cdot k^i$, where $w(u, v)$ is an N-dimensional vector of weights, $N = (n_1+3)(n_2+3)$, and k^i is an N-dimensional vector of coefficients corresponding to F_i, $i = 1, 2$. The b-spline regularity can be imposed with a discrete version of its thin-plate energy E_{TP}.

The linear systems proposed in Section 4 are over-determined, their first $2N$ unknowns being the coefficients of the two coordinates of the B-spline approximation of a smooth 2D function. They can be expressed in matrix form as $A \cdot X = B$. As a first option, we may take as solution the least squares minimum of:

$$E_{L_2}(X) = \|A \cdot X - B\|^2 \; + \; \lambda \, E_{TP}(k^1, k^2) \; , \tag{18}$$

for some constant $\lambda \geq 0$ (in practice only used for optical flow computation). In order to make the estimation less sensitive to outliers, and to possibly big (non-Gaussian) errors, we can use a smoothed L^1 penalization for the errors:

$$E_{L_{1\varepsilon}}(X) = \sum_{i=1}^{n} \sqrt{(A^i \cdot X - B_i)^2 + \varepsilon^2} \; + \; \lambda \, E_{TP}(k^1, k^2) \; , \tag{19}$$

where A^i is the i-th row of A, B_i is the i-th component of B, and $\varepsilon > 0$ is a small constant (in practice, $\varepsilon = 10^{-3}$). We can minimize $E_{L_{1\varepsilon}}$ via iteratively reweighted least squares. The SOCP problems arising when imposing the positive definiteness on G_ω in (15) can be solved using branch and bound [19].

References

1. Sturm, P., Ramalingam, S., Tardif, J.P., Gasparini, S., Barreto, J.: Camera models and fundamental concepts used in geometric computer vision. Foundations and Trends in Computer Graphics and Vision 6, 1–183 (2010)
2. Grossberg, M., Nayar, S.: A general imaging model and a method for finding its parameters. In: Proc. ICCV. IEEE (2001)
3. Pless, R.: Using many cameras as one. In: Proc. CVPR. IEEE (2003)
4. Sturm, P., Ramalingam, S.: A generic concept for camera calibration. In: Pajdla, T., Matas, J(G.) (eds.) ECCV 2004. LNCS, vol. 3022, pp. 1–13. Springer, Heidelberg (2004)
5. Diez-Cañas, G., Vasilyev, Y., Adato, Y., Zickler, T., Gortler, S., Ben-Shahar, O.: A linear formulation of shape from specular flow. In: Proc. ICCV. IEEE (2009)
6. Vasilyev, Y., Zickler, T., Gortler, S., Ben-Shahar, O.: Shape from specular flow: is one flow enough? In: Proc. CVPR. IEEE (2011)
7. Hartley, R., Kang, S.: Parameter-free radial distortion correction with centre of distortion estimation. In: Proc. ICCV. IEEE (2005)
8. Grompone von Gioi, R., Monasse, P., Morel, J.M., Tang, Z.: Towards high-precission lens distortion correction. In: Proc. ICIP. IEEE (2010)
9. Tardif, J.-P., Sturm, P., Roy, S.: Self-calibration of a general radially symmetric distortion model. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3954, pp. 186–199. Springer, Heidelberg (2006)
10. Ramalingam, S., Sturm, P., Boyer, E.: A factorization based self-calibration for radially symmetric cameras. In: Proc. 3DPVT (2006)
11. Thirthala, S., Pollefeys, M.: Radial multi-focal tensors. Int. J. Comput. Vis. 96, 195–211 (2011)
12. Ramalingam, S., Sturm, P., Lodha, S.: Towards generic self-calibration of central cameras. In: Proc. OMNIVIS. INRIA (2005)
13. Ramalingam, S., Sturm, P., Lodha, S.K.: Generic self-calibration of central cameras. Comput. Vis. Image Understand. 114, 210–219 (2010)
14. Taddei, P., Espuny, F., Caglioti, V.: Planar motion estimation and linear ground plane rectification using an uncalibrated generic camera. Int. J. Comput. Vis. 96, 162–174 (2012)
15. Grossman, E., Gaspar, J., Orabona, F.: Discrete camera calibration from pixel streams. Comput. Vis. Image Understand. 114, 198–209 (2010)
16. Nistér, D., Stewénius, H., Grossmann, E.: Non-parametric self-calibration. In: Proc. ICCV. IEEE (2005)
17. Espuny, F.: A closed-form solution for the generic self-calibration of central cameras from two rotational flows. In: Proc. VISAPP (2007)
18. Espuny, F., Burgos Gil, J.: Generic self-calibration of central cameras from two rotational flows. Int. J. Comput. Vis. 91, 131–145 (2011)
19. Sturm, J.: Using SEDUMI 1.02, a Matlab toolbox for optimization over symmetric cones (Updated for version 1.05). Tilburg Univ. (2001)
20. Horn, B., Schunck, B.: Determining optical flow. Artif. Intell. 17, 185–203 (1981)
21. Bogdanova, I., Bresson, X., Thiran, J.P., Vandergheynst, P.: Scale space analysis and active contours for omnidirectional images. IEEE Trans. Pattern Anal. Mach. Intell. 16, 1888–1901 (2007)
22. Puig, L., Guerrero, J.J.: Scale space for central catadioptric systems: Towards a generic camera feature extractor. In: Proc. ICCV. IEEE (2011)
23. Dierckx, P.: Curve and Surface Fitting with Splines. Oxford Univ. Press (1993)