INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Calibration of a stereo-vision system by the non-linear optimization of the motion of a calibration object

Matthieu Personnaz — Peter Sturm

N° 0269

September 2002

THÈME 3

ISSN 0249-0803

INRIA
RHÔNE-ALPES

Calibration of a stereo-vision system by the non-linear optimization of the motion of a calibration object

Matthieu Personnaz , Peter Sturm

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet MOVI

Rapport technique n° 0269 — September 2002 — 15 pages

Abstract: A common and high-performance practice when carrying out euclidian reconstruction from two cameras is to first calibrate both cameras thanks to a pair of images of a calibration object. However, it remains difficult to accurately determine the intrinsic parameters of each camera and we know their value have a significant effect on Euclidian reconstruction.
Moreover, the calibration issued from a single pair of images of the calibration object do not exploit the device to the full. Indeed, measures used to perform the calibration may arbitrarily grow by the use of several pairs of images of the calibration object.
It is the purpose of this paper to provide a method allowing to take advantage of such measures in order to improve simultaneously the accuracy of the intrinsic parameters and those of the ultimate Euclidian reconstruction. The method is based on a statistical and non-linear estimation of the parameters of the stereo rig.

Key-words: stereo-vision, calibration, euclidian reconstruction, non-linear optimization of the euclidian motion

Unité de recherche INRIA Rhône-Alpes
655, avenue de l'Europe, 38330 Montbonnot-St-Martin (France)
Téléphone : +33 4 76 61 52 00 — Télécopie +33 4 76 61 52 52

Calibration d'une paire de cameras par optimisation non lineaire du mouvement d'une mire de calibration

Résumé : Une méthode usuelle et performante pour effectuer une mesure 3D à partir d'un système de vision stereo consiste, dans un premier temps, à calibrer le système à partir d'une paire d'images d'un objet particulier, appelé mire de calibration. Toutefois, l'estimation précise des paramètres intrasèques des deux caméras reste difficile. Et nous savons que leurs valeurs ont une influence sur la précision de la reconstruction euclidienne.
D'autre part, une calibration déduite d'une unique paire d'images de la mire n'exploite pas l'ensemble des données auxquelles le dispositif permet d'accéder. En effet, le nombre de mesures peut être arbitrairement augmenter en prenant plusieurs paires d'images de la mire.
L'obectif de ce document est de fournir une méthode exploitant de telles mesures en vue d'augmenter simultanément la précision des paramètres internes et celle de la reconstruction euclidienne. Cette méthode est basée sur une estimation statisque et non-linéaire des paramétres du système stereo.

Mots-clés : stereo-vision, calibration, reconstruction euclidienne, optimisation non-linéaire du mouvement euclidien

1 Introduction

The quality of the calibration of a stereo-vision rig determines partially the accuracy of euclidian reconstruction ([1]). The robust and high-performance calibration methods usually require both the 3D coordinates of some points of the viewed scene and the measures of the 2D coordinates of their images ([3]). The calibration object refers to a specific object which contains points whose 3D coordinates are known.
However, given the previous device, the estimation of the intrinsic parameters of each camera from a pair of images of the calibration object remains difficult ([10]). So this approach does not allow to know if the estimated parameters will permit a reconstruction with an optimal accuracy.
Now, an ideal solution could be the use of a calibration object whose position and number of known points would be arbitrarily fixed. Indeed, assuming the parameters are maximum likelihood estimates, the likelihood increases with the number of measures ([11]). Such a device is not technically feasible. On the other hand, to move the calibration object and gather pairs of images of the object at distinct positions allows to increase arbitrarily the number of measures.
In order to exploit, according to a maximum likelihood way, the set of measures issued from the last process, this paper defines a non-linear optimization of both the camera parameters and the successive positions, in a camera reference, of the calibration object.
First, this paper defines a model of the stereo rig and the motion of the calibration object. A parameterization of the model is provided as well.
Then, an initial solution is determined thanks to a statistical approach. This solution is used as data input by a Levenberg-Marquardt algorithm. This paper describes such an iterative algorithm, taking advantage of the sparse structure of the Jacobian matrices.
At last, experimental results illustrate and quantify the contributions of the model and the optimization both on the value of the maximum likelihood and the accuracy of the euclidean reconstruction.

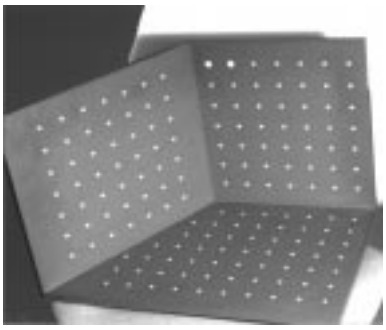

Figure 1: The calibration object

2 Model of the calibration device

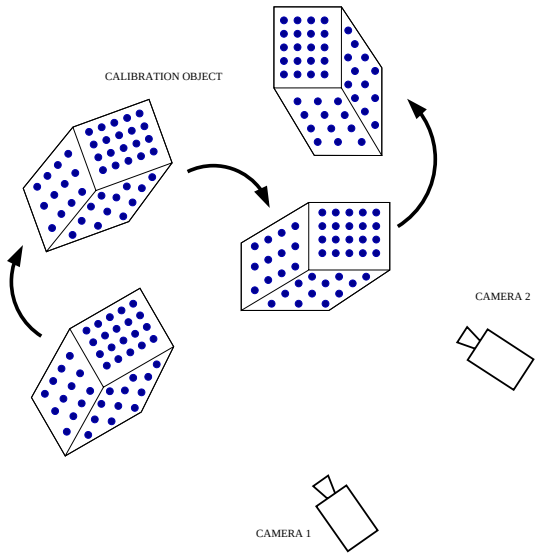

Figure 2: Motion of the calibration pattern in front of the two cameras

2.1 Model of the stereo-vision rig

Each camera is modelled as a finite projective camera: the calibration matrix of each one has the form

$$K = \begin{pmatrix} \alpha_x & s & x_0 \\ 0 & \alpha_y & y_0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1}$$

The value of the skew parameter s is also not neglected.
For added generality, the correction of the radial distortion is carried out according to the model defined in ([7]). In pixel coordinates, the model has the form

$$\begin{pmatrix} x_d - x_c \\ y_d - y_c \end{pmatrix} = L(\tilde{r}) \begin{pmatrix} \tilde{x} - x_c \\ \tilde{y} - y_c \end{pmatrix} \tag{2}$$

where

- $(\tilde{x}, \tilde{y})$ is the image position which obeys linear projection
- (x_d, y_d) is the actual image position submit to radial distortion

- (x_c, y_c) is the center of radial distortion
- $\tilde{r} = \sqrt{(\tilde{x} - x_c)^2 + (\tilde{y} - y_c)^2}$ and $L(\tilde{r}) = 1 + k_1\tilde{r} + k_2\tilde{r}^2$

The relative position of the both cameras is represented through the transformation of the first to the second camera reference. If a world point has respectively X and X' as vectors of coordinates in the coordinate systems associated to the first and the second camera, the transformation is written

$$X' = RX + t \tag{3}$$

where R and t are the rotation and the translation defining the transformation.
Now, if K and K' denote the respective calibration matrices of the first and the second camera, the stereo rig can be modelled by the pair of perspective matrices $P = K[I|0]$ and $P' = K'[R|t]$, and by the model of distortion of each camera.
So, to parameterize the model of the stereo rig, we choose the non-zero coefficients of the K and K' matrices, the coordinates of a unit quaternion representing R and those of the translation t. The radial distortion of each camera is parameterized by the coefficients k_1, k_2, and the coordinates (x_c, y_c) of the center of distortion.

2.2 Model of the motion of the calibration object

We assume in this paper that a reference is attached to the calibration object. The successive positions of the calibration object are also related to the coordinate system of the first camera by the transformation of the object reference to the first camera reference. If $[R_j|t_j]$ is the jth transformation associated to the jth position of the calibration object, the transformation formula is written

$$Y = R_jX + t_j \tag{4}$$

where Y and X are respectively the coordinates of a world point in the camera reference and in the object reference. The chosen parameterization is made up of the set of quaternions $(q_j)_j$ representing the rotation $(R_j)_j$, and the set of translations $(t_j)_j$.

3 Estimation of the model

The purpose of this section is to define a merit function χ and to provide an algorithm which determines, by the minimization of χ, best-fit parameters of our model.
The minimization of the cost function is studied over two sets of parameters: the set of the stereo rig parameters and the set of the parameters associated to both the motion of the calibration object and the stereo rig. For the first minimization, a Levenberg-Marquard is implemented. The second minimization algorithm is as well a Levenberg-Marquard algorithm; this part focus on it in order to reduce the computation time. The section defines as well an initial solution used as data input of the previous iterative algorihms.

3.1 Expression of the cost function

By assuming that the image coordinate measurement errors obey a centered-normal distribution, the Maximum Likelihood Estimate (MLE) of the motion of the calibration object and the stereo rig minimizes the geometric error ([4]), i.e.:

$$\chi = \sum_l \|x_l - \hat{x}_l\|^2 \tag{5}$$

where $(x_l)_l$ is the set of vectors associated to the measured 2D image points, $(\hat{x}_l)$ is the set of the estimated 2D points and $\|.\|$ is the euclidean norm.
Now, to express χ according to the parameters of our model it is enough to provide a convenient expression of the images of a world point viewed by the two cameras. In the following part, all entities are represented in homogeneous coordinates, i.e. are defined up to scale. The symbol $\sim$ means "equals up to scale". To simplify the formulas, the effects of the radial distortion are omitted as well.
Let $(X_i)_{1\leq i\leq n}$ be the set of vectors representing the known coordinates, in a reference attached to the calibration object, of world points lying on the calibration object.
Let x_{ij} and x'_{ij} $(1 \leq i \leq n, 1 \leq j \leq m)$ be the images respectively in the first and the second camera of the point X_i at the jth position of the calibration object.
Let $[P_j, P'_j]$ $(1 \leq j \leq m)$ be a conjugated pair of perspective matrices associated to the jth position of the calibration object. Such a pair is defined by the system

$$\left\{ \begin{array}{lcll} x_{ij} & \sim & P_j X_i & (1 \leq i \leq n, 1 \leq j \leq m) \\ x'_{ij} & \sim & P'_j X_i & \end{array} \right. \tag{6}$$

Now, if we denote $(H_j)_{1\leq j\leq m}$ the set of homographies corresponding to the transformations $([R_j|t_j])_{(1\leq j\leq m)}$ (see section 2.2 on page 5), i.e.

$$H_j \sim \left(\begin{array}{cc} R_j & t_j \\ 0 & 1 \end{array} \right) \; 1 \leq j \leq m \tag{7}$$

we obtain the system

$$\left\{ \begin{array}{lcll} x_{ij} & \sim & P H_j X_i & (1 \leq i \leq n, 1 \leq j \leq m) \\ x'_{ij} & \sim & P' H_j X_i & \end{array} \right. \tag{8}$$

This is issued from the equations $P_j \sim P H_j$ and $P'_j \sim P' H_j$ $(1 \leq j \leq m)$.

3.2 Estimation of the initial solution

3.2.1 Initialization of the parameters of the stereo rig

For each position of the calibration object, the perpective matrices of the two cameras, i.e. the set $([P_j, P'_j])_{1\leq j\leq m}$ and the associated parameters of distortion are computed thanks to

an algorithm adapted from [3].
To estimate the intrinsic parameters, the calibration matrices K_j and K'_j are first extracted thanks to the formula defined in [6]. The distribution of the intrinsics parameters being roughly gaussian ([10]), the coefficients of the camera matrices K and K' are estimated by the means of the coefficients repectively of the calibration matrices $(K_j)_{1\leq j\leq m}$ and $(K'_j)_{1\leq j\leq m}$. The method is extended to the coefficients of distortion.
From each conjugated pair $[P_j, P'_j]$ $(1 \leq j \leq m)$, an estimation $[Q_j|s_j]$ of the relative position $[R|t]$ of the two cameras is deduced. To take account of the m estimations of $[R|t]$, the translation t is estimated by the mean of the m translations $(s_j)_{1\leq j\leq m}$. The estimation of the rotation R is performed by applying the algorithm described in [9] to the m rotations $(Q_j)_{1\leq j\leq m}$

$$\bar{Q}_{t+1} = \bar{Q}_t \frac{1}{m} \sum_{j=1}^{m} (\bar{Q}_t^{-1} Q_j) \tag{9}$$

where $\bar{Q}_t$ represents the estimation of the rotation R at the step t. The process is stopped when $\sum_{j=1}^{m} \|\bar{Q_j}^T Q_j\| < 10^{-10}$.

3.2.2 Initialization of the motion

At the jth $(1 < j < m)$ position of the calibration object, some 3D points of the calibration object which are simultaneouly viewed by both cameras are reconstructed by performing the triangulation algorithm defined in [2]. The coordinates of each 3D point are also computed in the first camera reference by taking as input the conjugated perspective matrices $[P_j, P'_j]$, the associated coefficients of distortion and the coordinates of the matched image points.
To this set of these space points reconstructed in the first camera reference corresponds the given set of their 3D coordinates in the calibration object reference. The transformation $[R_j, t_j]$ which matches the both set of 3D is deduced by applying the algorithm described in [12].
The motion is then initialized by estimating the set $([R_j, t_j])_{1\leq j\leq m}$.

3.3 Minimization over the motion and the stereo rig

To minimize the χ function over the parameters of the motion and those of the stereo rig, a Newton-type iterative method had been implemented. More precisely, the used algorithm applies the framework of the sparse Levenberg-Marquardt algorithm defined in [5], p.577. This algorithm is adapted to our minimization problem since one observes that the measurements of the 2D points in a conjugated pair of images at the jth position of the calibration object $(1 \leq j \leq m)$ do not depend on the kth position of the calibration object if $k \neq j$ $(1 \leq k \leq m)$. This part describes how to exploit the framework. We denote $\hat{z}$ the estimated value of any measured variable z and v^T the transpose of any vector v.

Let the parameter vector partition as $P = (a^T, b_1^T, ..., b_m^T)$ with

$$\begin{array}{rcl} b_j^T & = & (q_j^T, t_j^T) \ (1 \leq j \leq m) \\ a^T & = & (h^T, h'^T) \end{array} \tag{10}$$

where the vector h and h' are respectively made up of the parameters associated to each camera, i.e. the calibration matrices K and K' (see section 2.1 on page 4) and the coefficients of distortions. The vector h' contains as well the coefficients of the transformation $[R, t]$ defining the relative position of the cameras. The transformations $(q_j, t_j)_j$ has been previously defined in the part 2.2, page 5.
Let $X^T = X_1^T, \ldots, X_m^T$ be the measurement vector. X_j $(1 \leq j \leq m)$ corresponds to the image points of some known space points of the calibration object viewed by the two cameras, at the jth position of the calibration object. For instance, assume that $(x_{ij})_{1 \leq i \leq s}$ and $(x'_{ij})_{1 \leq i \leq s}$ are the measured image points, respectively in the first and the second camera, of some known space points $(M_i)_{1 \leq i \leq s}$ lying on the calibration object. Furthermore, $(x_{ij})_{1 \leq i \leq s}$ and $(x'_{ij})_{1 \leq i \leq s}$ are associated to the jth position $(1 \leq j \leq m)$. Then, one writes X_j as

$$X_j^T = x_{1j}^T, \ldots, x_{lj}^T, x_{1j}^{'T}, \ldots, x_{lj}^{'T} \tag{11}$$

With such notations, note that the independence between measures and positions of the calibration object is written

$$\frac{\partial \hat{X}_i}{\partial b_j} = 0 \text{ if } i \neq j \tag{12}$$

Let us now focus on the computation of the jacobian matrix

$$J = \left(\begin{array}{c|ccc} A_1 & B_1 & & \\ A_2 & & B_2 & \\ \vdots & & & \ddots \\ A_m & & & B_m \end{array} \right) \tag{13}$$

with

$$A_j = \frac{\partial \hat{X}_j}{\partial a} \tag{14}$$

$$B_j = \frac{\partial \hat{X}_j}{\partial b_j} \tag{15}$$

The Jacobian matrices $(A_j)_j$ takes in our case the special form

$$A_j = \left(\begin{array}{cc} \frac{\partial \hat{x}'_{ij}}{\partial h} & 0 \\ 0 & \frac{\partial \hat{x}'_{ij}}{\partial h'} \end{array} \right) \tag{16}$$

Indeed, from the system 8, we deduced

$$\frac{\partial \hat{x}'_{ij}}{\partial h} = 0 \tag{17}$$

$$\frac{\partial \hat{x}_{ij}}{\partial h'} = 0 \tag{18}$$

The design of the minimization algorithm is then achieved by implementing the five previous formula in the framework described in [5], p577.

4 Experiments

4.1 Goal

The goal of the experiments is to quantify the quality of the stereo rig calibration methods. The quality of the calibration is mainly assessed using two values:

- the reconstruction accuracy.
- the reprojection error, whose computation from 8 is supposed to be known.

Three methods are also compared:

- ONE corresponds to the standard calibration from a single pair of images of the calibration object.
- INITIAL performs the calibration from several pairs of images. The retained parameters are the initial solution defined in part 3.2.2.
- LM uses a Levenberg-Marquardt algorithm to minimize the cost function χ over the stereo rig parameters.
- SLM applies the sparse Levenberg-Marquardt algorithm described in 3.3 to minimize χ over the stereo rig and the motion.

4.2 Experimental process

The stereo rig is made up of two JAI cameras distant from $1 \pm 0.1m$. The cameras point to the calibration object. For all the positions of the calibration object, the distance between a camera and the calibration object is contained in the interval $2.5 \pm 0.5m$.
Then, the software *tele2-3.2* (see [10]) is used to gather stereo pairs of the calibration object and to compute two kinds of data:

- A set of several lists of pairs of matched image points. Each list, whose size may possibly vary, corresponds to a position of the calibration object. Each element of a list is a pair of 2D points, images in each camera of a known 3D world point lying on the calibration object. At last, the positions of the calibration object are mutually distinct.

- The set $([P_j, P'_j])_j$ of the conjugated perpespective matrices and the associated coefficients of distortion defined in the section 3.2.1 page 6.

Let now precise in which processes these data are used. To each position of the calibration object is associated a pair of images of the calibration object. And, from each pair of images, a set of pairs of matched image points is extracted; this set is referred to as a set of 2D-2D matched points. A set of 2D-2D matched points may be used either to calibrate the device or to build 3D points by assuming that calibration results are provided. So, we have partitionned the set of pairs of images of the calibration object into two subsets. The first one is used to calibrate the stereo rig through the methods addressed in the paper. The second subset is dedicated to the estimation of the reconstruction accuracy corresponding to each calibration. Subsequently the 2D measured data used for the calibration are different from those used to estimate the reconstruction accuracy.
To achieve the experimental process, let us describe the algorithm allowing to estimate the reconstruction accuracy. From a pair of 2D matched points and the calibration results, the associated world point is reconstructed in the first camera reference by applying a triangulation algorithm based on [2]. By reiterating the process for the available pairs of matched points, a list of 3D world points $(\tilde{Q}_l)_{1\leq l\leq s}$ is also reconstructed in the first camera reference. These points are supposed to coincide with some known points of the calibration object whose coordinates $(Q_l)_{1\leq l\leq s}$, in the calibration object reference, are given. The transformation between the first camera reference and the calibration object reference is computed by applying the fitting algorithm defined in [12] to sets of 3D points. The found transformation also maps the coordinates of the reconstructed points, in the first camera reference, into the calibration object reference. Once the reconstructed points are expressed in the calibration object reference, the accuracy δ is estimated by the mean of the euclidean distances between the reconstructed world points and the given world points:

$$\delta = \frac{1}{s}\sum_{l=1}^{l=s} \|\tilde{Q}_l - Q_l\| \tag{19}$$

The standard deviation σ is estimated from the usual formula:

$$\sigma^2 = \frac{1}{s(s-1)}\sum_{l=1}^{s} (\|\tilde{Q}_l - Q_l\| - \delta)^2 \tag{20}$$

4.3 Results

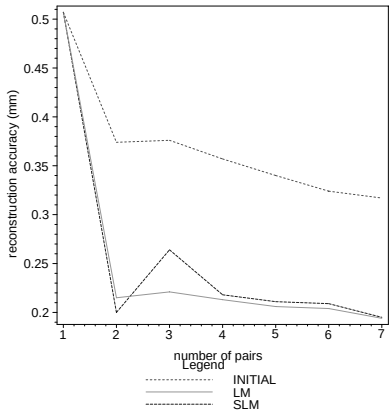

Figure 3: Comparison of the methods INITIAL, SLM and LM: the average of the accuracy reconstruction related to the number of positions of the calibration object. The number of reconstructed points is 755. The reconstrution method is defined in section 4.2.

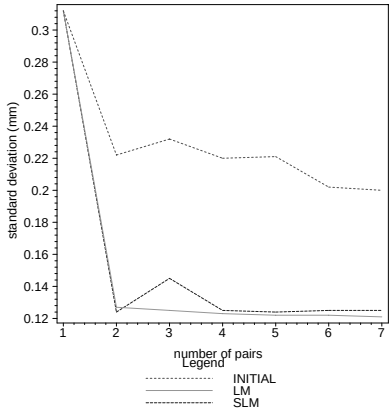

Figure 4: Comparison of the methods INITIAL, SLM and LM: the standard deviation of the accuracy reconstruction related to the number of positions of the calibration object. The number of reconstructed points is 755. The reconstrution method is defined in section 4.2.

Figures 1 and 2 shows the improvement of the reconstruction accuracy when the number of measures grows. The reconstruction error also decreases by 40% for the INITIAL method: 0.31mm for INITIAL, for 7 pairs of images, versus 0.51mm for ONE. For the other two

iterative methods, the diminution of the reconstuction error is more pronounced. Indeed, we notice it allows an additional decrease by approximatively 30%: 0.19mm for SLM versus 0.31mm for INITIAL for 7 pairs of images.
The gain provided by the use of an iterative algoritm is as well noteworthy: the decrease of the reconstruction error exceeds generally 30% by using LM or SLM versus INITAL.
Nevertheless, from 4 and more pairs of images, the SLM method gives slighly worse results than the LM method. So, in our case, there is no advantage to minimize the merit function χ over the motion. That means the estimation of the motion by the initial solution is better than the estimation deduced from SLM. This last result may be explained by estimates corresponding to a local minimum of the merit function.

The following table shows the evolution of the averages and the standard deviations of the reprojection errors (in pixel) for the different optimization processes.

Number	1		2		3		4		5		6		7	
of pairs	ave.	std.	ave.	std.	ave.	std.	ave.	std.	ave.	std.	ave.	std.	ave.	std.
INITIAL	.63	.48	.37	.29	.59	.33	.56	.26	.51	.26	.46	.26	.44	.28
LM	.63	.48	.10	.07	.37	.16	.40	.14	.38	.16	.36	.18	.34	.2
SLM	.63	.48	.089	.071	.36	.15	.36	.16	.32	.17	.30	.18	.30	.17

The table of the reprojection errors illustrates the decrease of the cost function for the two optimization methods LM and SLM. The decrease provided by both optimization processes versus the INITIAL method is between 20% and 30%. The gain privided by SLM compared to LM is quite low: roughly 15%. As we have shown previously, this last slight decrease does not allow an improvement of the reconstruction accuracy.

5 Conclusion

This paper has studied the process of calibration from several positions of the calibration object. It has shown the improvement of the reconstruction accuracy just from a statistical approach, i.e. without optimizing the stereo rig nor the motion of the calibration object. The influence of the accuracy of the intrinsic parameters on the reconstruction error had also been illustrated.
However we have illustrated the advantage of a non-linear minimization of the parameters. More precisely, to estimate the mean of a parameter with a convenient precision requires a minimum number of measures, say 30 (see [10]). In our experiment we use only 7 positions of the calibration object, and then 7 measures of the stereo rig parameters to estimate their mean. This may obviously be not sufficient and this low number is a consequence of the time taken by the calibration process. Indeed, using the software *tele2-3.2* (see [10]), the time to interactively gather a stereo pair of images, extract and match 2D image points is roughly 1mn. The time required by the minimization algorithm is neglected: less than 1s. That's why the number of pairs has been limited in our experiment to 7. In such a case, the

optimization processes may then possibly used to compensate this lake of measures.
At last, the calibration and reconstructed results are based on the great accuracy of the given 3D world points and the measures of their 2D images. A futur work could be to study the problem when the error on the 3D given world point can't be neglected ([8]) or when their 2D images are computed with a low precision.

Contents

References

[1] R.I. Hartley and R. Kaucic. Sensitivity of calibration to principal point position. In *Computer Vision-ECCV*, volume 2, pages 433–446, 2002.

[2] R.I. Hartley and P. Sturm. Triangulation. *Computer Vision and Understanding*, 68(2):146–157, 1997.

[3] R.I. Hartley and A. Zisserman. Camera estimation from a calibration object. In *Multiple View Geometry in Computer Vision*, page 170. Cambridge University Press, 2001.

[4] R.I. Hartley and A. Zisserman. Geometric error. In *Multiple View Geometry in Computer Vision*, page 169. Cambridge University Press, 2001.

[5] R.I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2001.

[6] R.I. Hartley and A. Zisserman. The projective camera. In *Multiple View Geometry in Computer Vision*, page 150. Cambridge University Press, 2001.

[7] R.I. Hartley and A. Zisserman. Radial distortion. In *Multiple View Geometry in Computer Vision*, pages 179–180. Cambridge University Press, 2001.

[8] Marc Viala J.M. Lavest and Michel Dhome. Do we really need an accurate calibration pattern to achieve a reliable camera calibration? In *Computer Vision-ECCV*, volume 1, pages 158–174, 1998.

[9] X. Pennec. Computing the mean of geometric features, application to the mean rotation. *Inria Research Report*, 3371:8, 1998.

[10] M. Personnaz and R. Horaud. Camera calibration: estimation, validation and software. *Inria Technical Report*, 258:11–19, 2002.

[11] G. Saporta. La méthode du maximum de vraisemblance. In *Probabilités, analyse des données et statistique*, pages 301–302. TECHNIP, 1990.

[12] J. Tarel. Recalage géométrique avec plusieurs prototypes. *Inria Research Report*, 2988:11, 1996.

Unité de recherche INRIA Rhône-Alpes
655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)
Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-0803