



Combining Off- and On-line Calibration of a Digital Camera

Magdalena Urbanek, Radu Horaud, Peter Sturm

► To cite this version:

Magdalena Urbanek, Radu Horaud, Peter Sturm. Combining Off- and On-line Calibration of a Digital Camera. 3rd International Conference on 3D Digital Imaging and Modeling, May 2001, Quebec City, Canada. pp.99-106, 10.1109/IM.2001.924413 . inria-00525659

HAL Id: inria-00525659

<https://inria.hal.science/inria-00525659>

Submitted on 26 May 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Combining Off- and On-line Calibration of a Digital Camera

Magdalena Urbanek, Radu Horaud and Peter Sturm

INRIA Rhône-Alpes

655, avenue de l'Europe, 38334 Saint Ismier Cedex, France
{Magdalena.Urbanek, Radu.Horaud, Peter.Sturm}@inrialpes.fr

Abstract

We introduce a novel outlook on the self-calibration task, by considering images taken by a camera in motion, allowing for zooming and focusing. Apart from the complex relationship between the lens control settings and the intrinsic camera parameters, a prior off-line calibration allows to neglect the setting of focus, and to fix the principal point and aspect ratio throughout distinct views. Thus, the calibration matrix is dependent only on the zoom position. Given a fully calibrated reference view, one has only one parameter to estimate for any other view of the same scene, in order to calibrate it and to be able to perform metric reconstructions. We provide a close-form solution, and validate the reliability of the algorithm with experiments on real images. An important advantage of our method is a reduced - to one - number of critical camera configurations, associated with it. Moreover, we propose a method for computing the epipolar geometry of two views, taken from different positions and with different (spatial) resolutions; the idea is to take an appropriate third view, that is "easy" to match with the other two.

1. Introduction

The problem of recovering the Euclidean structure of a scene is strongly associated with the estimation of the camera internal parameters, i.e. calibration. When no calibration knowledge provided, one can reconstruct only a projective model of the scene [6, 10].

1.1. Previous work

The most basic solution to compute the internal parameters employs a calibration grid or planes, and performs an off-line calibration. However the restriction of keeping an identical camera state (including zooming and focusing) while shooting subsequent images can hardly be fulfilled in practice.

Another idea is to self-calibrate an entire sequence. Existing approaches follow several directions. One is to assume invariance of unknown intrinsic parameters throughout distinct views [12, 15, 1, 8], thus not to allow for zooming/focusing, which is quite a strong constraint. Given a stereo pair of an arbitrary scene, one cannot vary but *the magnification parameter* (we use that term, to avoid confusion of associating different meanings to *the focal length*, in vision and optics), while having the other ones known [7]. Other methods [3, 13] allow the retrieval of varying magnification parameter and fixed principal point. Furthermore, if provided with at least 9 views, it is possible to fix only one camera internal parameter and let the other ones vary [13, 11].

In reality, such a general calibration problem cannot be solved reliably. On the other hand, one can quite easily provide some prior information, which simplifies the task. Our approach belongs to such a group of techniques.

1.2. Motivation

All considered cases of self-calibration, which allow magnification parameter variation throughout distinct views, suffer from a significant number of critical camera configurations [14]. It is therefore much "safer" not to change the camera settings.

Let us combine one fully calibrated image (*the reference image*) with an uncalibrated one, taken from a different viewpoint. Then, one has only one magnification parameter to estimate. What about the other intrinsic parameters? The complex relationship between calibration and camera lens control settings [16] does not allow straight-forward simplifications.

To summarize, we are interested in the following issues:

- Are there any conditions that enable the use of a priori knowledge of the intrinsic parameters?
- Can one allow for zooming/focusing, while still maintain a small family of critical situations?
- What can be done with stereo pairs, if one camera/view is fully calibrated?

1.3. Contribution

We combine off-line and on-line methods in order to calibrate a digital camera with a zoom lens and auto-focus.

We introduce a novel outlook on the self-calibration problem, by reducing to one the number of intrinsic parameters to be estimated. We provide a close-form solution for the method. Also, one has to account for only a single family of critical camera configurations [14].

By studying the behaviour of the camera intrinsic parameters as a function of variable zoom and focus, we derive approximate values for the aspect ratio and the principal point. We identify a small influence of focus upon calibration, which becomes negligible for settings larger than 2.5m. We conclude, that once a camera is calibrated for a known zoom setting, one can re-use those values any time that zoom is set. Therefore, we recommend employing minimally or maximally zoomed-in images as the reference ones, since those zoom settings can be reliably reproduced.

Furthermore, we simplify the computation of the epipolar geometry for stereo images of different resolutions, omitting a direct matching between them. The problem of matching two images of different zoom and viewpoint is therefore decomposed into two simpler matching problems: a wide baseline matching with the same zoom [2], and matching images with different zoom, shot from the same viewpoint [5].

The proposed method of "combined calibration" estimates the intrinsic parameters with even 2%-accuracy, from real images, leading to a reliable Euclidean reconstruction.

2. Camera modeling

2.1. The model

We assume the perspective camera model with the projection matrix of the form:

$$P = K \begin{pmatrix} R & t \end{pmatrix} \quad (1)$$

where R and t represent the orientation and the position of the camera with respect to the world coordinate system, and K is the calibration matrix:

$$K = \begin{pmatrix} k\alpha & 0 & u_0 \\ 0 & \alpha & v_0 \\ 0 & 0 & 1 \end{pmatrix}$$

with the principal point (u_0, v_0) , the magnification parameter α and the aspect ratio k . We assume a zero-skew.

A scene point M is projected onto the image onto a point m via $m = PM$.

2.2. Off-the-shelf digital camera

Most often one is provided with digital cameras, which allow mechanical setting of both zoom and focus. One can specify the area of interest (and thus, its depth on the image) and focus on chosen features within the area.

We have worked with the Olympus Camedia C-2500L digital camera. It provides both auto-focus and manual-focus with discretized values from 0.3m until 15m and ∞ to be set. The zoom, on the contrary, has a continuous range and a manual drive, which makes the reproducibility of different settings difficult (with notable exceptions for the minimal and the maximal zooms).

Each $(zoom, focus)$ setting corresponds to a physical configuration of lenses, inside the camera. Since their functional dependencies are complex, we cannot specify the exact camera state, which makes the estimation of camera internal parameters difficult. When using auto-focus, the only camera settings that we are able to reproduce (and to expect the same calibration results, for an arbitrary image, taken with the same settings) are: $(zoom, focus) = (zoom_{min}, \infty)$ and $(zoom, focus) = (zoom_{max}, \infty)$.

The question is how do the entries of the calibration matrix K change with variations of zoom and focus. Experiments described in the following section suggest conditions, under which the internal camera calibration can be assumed invariant, for different $(zoom, focus)$ settings.

3. Off-line stability study of calibration

We study the stability of the camera internal parameters, under change in the camera mechanical settings, zoom and focus. We point out the parameters that do not vary much, and can be assumed invariant. We find a small influence of focus on calibration, if the camera is far enough from the scene. Finally, we provide calibration knowledge for particular zoom settings, which is to be used a priori, in self-calibration.

3.1. A way to calibrate

We extract the calibration matrix K from the projection matrix P , estimated from correspondences between non-coplanar 3D points and their 2D images.

The form of $P = \begin{pmatrix} \bar{P} & p \end{pmatrix}$ and (1) imply: $\bar{P} = KR$. Since $\bar{P}\bar{P}^T = KRR^TK^T = KK^T$, we can simply obtain K from the Cholesky decomposition of $\bar{P}\bar{P}^T$.

In order to estimate P , we run a non-linear algorithm, which minimizes the reprojection error

$$C = \sum_{i=1}^n (u_i - u_{mi})^2 + (v_i - v_{mi})^2 \quad (2)$$

of n image points (u_i, v_i) and reprojections (u_{mi}, v_{mi}) of the corresponding 3D points M_i .

3.2. Optical distortion

Since imperfect camera lenses give rise to non-perspective image distortion, it is often necessary to optimize (2) using additional distortion parameters. In some cases, this extended projection model causes over-parameterization, resulting in instabilities in the estimation of all intrinsic parameters.

Based on the observation, that the bigger the zoom used, the less distortion is present in the image, we can point out experimentally a "critical" zoom, for which the *estimated* distortion coefficient does not decrease with the increase of zoom. Therefore, we omit the distortion parameters in the optimization, if a zoom is bigger than the "critical" one.

We only estimate the first term D_r of the radial distortion, which proved sufficient to provide reliable results. Overall, the employed calibration method is described in [4].

3.3. Experiments

We stepped the lens through the full range of focus, while the zoom was examined in two positions: the minimal and the maximal ones. At each step, we performed a full camera calibration (images of a calibration grid were considered). To ensure the stability of calibration, we considered only images with a sufficiently large number of control points clearly visible.

We used manual focusing. For each (*zoom*, *focus*) setting, we took several images with slightly different orientations of the calibration grid. The distance camera-grid was kept identical to the value of the set focus.

We considered focus values between 1m and 5m. The images were of size 640×512 pixels. The obtained estimates of the internal camera parameters are listed separately: for the minimal zoom (Table 1), for the maximal zoom (Table 2).

3.4. Dependencies

What information can be extracted from Tables 1 and 2?

Aspect ratio (k). It is close to unity. The equality $k = 1$ is valid for any (*zoom*, *focus*) setting, with a relative error smaller than 0.2%.

Magnification parameter (α). For the minimal zoom, α stays constant relative to focusing. For the maximal zoom, the same is observed as soon as the distance camera-object is bigger than 2.5m (see Figure 1). Hence, for a chosen zoom, it is possible to represent the relevant α with a single value (e.g. the median of the estimates): $\alpha_{min} = 706$ (with 2%-relative error) for (*zoom*, *focus*) =

Focus[m]	$k[1]$	$\alpha[\text{pix}]$	$u_0[\text{pix}]$	$v_0[\text{pix}]$	$D_r[1]$
1	0.9993	700	321	268	-0.2393
	0.9991	698	321	267	-0.2423
	0.9999	700	317	267	-0.2363
1.2	0.9992	695	314	277	-0.2598
	0.9996	702	320	269	-0.2405
	0.9997	728	294	238	-0.1468
1.5	0.9998	731	316	232	-0.1574
	1.0007	710	325	269	-0.2469
	0.9998	723	318	234	-0.1601
2	1.0007	736	295	269	-0.1523
	1.0002	699	319	274	-0.2970
2.5	1.0001	722	318	268	-0.2207


Table 1. Calibration results: the *minimal* zoom and varying focus.

Focus[m]	$k[1]$	$\alpha[\text{pix}]$	$u_0[\text{pix}]$	$v_0[\text{pix}]$	$D_r[1]$
1	0.9996	921	316	268	-0.0976
	0.9994	920	318	268	-0.0945
	0.9992	918	319	269	-0.1037
1.2	1.0010	1133	320	266	0.0445
	1.0009	1122	317	274	0.0075
	1.0008	1128	318	271	0.0218
1.5	1.0015	1384	310	297	0
	1.0005	1386	291	296	0.0294
	1.0013	1391	320	292	0
2	1.0020	1749	312	313	0
	1.0020	1740	311	310	0
	1.0013	1745	289	303	0
2.5	1.0030	1969	301	324	0.1470
	1.0008	1944	255	316	0.0306
3.5	1.0024	1959	314	356	0.0295
	1.0012	1965	290	336	0
5	1.0016	1999	301	344	0.0182


Table 2. Calibration results: the *maximal* zoom and varying focus.

($zoom_{min}$, $focus \geq 0.3\text{m}$), and $\alpha_{max} = 1965$ (with 1%-relative error) for (*zoom*, *focus*) = ($zoom_{max}$, $focus \geq 2.5\text{m}$).

Principal point (u_0, v_0). Overall, it concentrates near the image centre (see Figure 2). Since in general, the exact position of the principal point does not have a big impact upon the quality of reconstruction, it is possible to employ approximate statistical values, obtained from the Student's reliability test: $u_0 = 311 \pm 21$, $v_0 = 280 \pm 42$, with a factor of risk 0.1. Further on, we will use the approximation: $(u_0, v_0) = (311, 280)$.



(a) Minimal zoom: median $\alpha_{min} = 706$.



(b) Maximal zoom: median $\alpha_{max} = 1965$.

Figure 1. Possibilities to approximate the magnification parameter α with its median.




Figure 2. Principal point concentrates near the image centre.

Auto-focusing. For a fixed zoom, the setting of focus does not influence calibration significantly. We can use auto-focusing, and still be capable to employ calibration results for the examined zooms. We only have to keep in mind

the requirement concerning the maximal zoom: the distance camera-scene has to be larger than 2.5m.

3.5. Final results to be used in self-calibration

A view taken with the minimal/maximal zooming. We are provided with calibration matrices of reference: K_{min} for the minimal zoom case (for any focus value), and K_{max} for the maximal zoom case (for focus ≥ 2.5 m).

A view taken with an arbitrary (unknown) zooming. One is provided with the values of k and (u_0, v_0) . Hence, α remains the only calibration parameter to determine.

A summary is given in Table 3.

Zoom	Focus[m]	$k[1]$	$\alpha[\text{pix}]$	$u_0[\text{pix}]$	$v_0[\text{pix}]$
<i>min</i>	≥ 0.3	1	706	311	280
<i>max</i>	≥ 2.5	1	1965	311	280
?	?	1	?	311	280

Table 3. Results of off-line calibration (the Olympus Camedia C-2500L digital camera).

4. Self-calibration

We consider a stereo pair: a calibrated reference image and an image taken with an unknown zoom. (In practice, we obtain the calibration for the reference image simply by taking it using the minimal or the maximal zoom, and adopting the according intrinsic parameters, obtained by the off-line calibration.) We are thus provided with calibration matrices: K_{ref} , fully known, for the reference image, and K , defined up to unknown α , for the other image. Due to Kruppa's equations [9], we derive a close form solution for α . Also, we reveal stereo configurations, for which our self-calibration algorithm fails.

4.1. Kruppa's equations

Finding the matrix K associated with a camera is equivalent to finding the image ω of the absolute conic, taken by that camera. Since $\omega^{-1} \sim KK^T$, let us denote $C = KK^T$ for the camera to be calibrated, and $C_{ref} = K_{ref}K_{ref}^T$ for the camera, that took the reference image.

The link between images of the absolute conic and the epipolar geometry can be expressed as follows ([17]; F is the fundamental matrix of the stereo; e is the epipole on the image, taken by the uncalibrated camera):

$$FC_{ref}F^T \sim [e]_{\times} C [e]_{\times}^T \quad (3)$$

Having separated the known entries of matrix K (k, u_0, v_0) from the unknown one (α):

$$K = \begin{pmatrix} k & 0 & u_0 \\ 0 & 1 & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} = K_0 K_\alpha$$

and multiplying (3) from the left by K_0^T , and from the right by K_0 , we obtain:

$$K_0^T F K_{ref} K_{ref}^T F^T K_0 \sim K_0^T [e]_\times K_0 \underbrace{K_\alpha K_\alpha^T}_{C_\alpha} K_0^T [e]_\times^T K_0 \quad (4)$$

with C_α of the following form:

$$C_\alpha = \begin{pmatrix} \alpha^2 & 0 & 0 \\ 0 & \alpha^2 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (5)$$

Let us denote:

$$\underline{E} = K_0^T F K_{ref} \quad (6)$$

which moves F to a "semi-calibrated" space. From a property (valid for any matrix A and vector \mathbf{v})

$$A^{-T} [\mathbf{v}]_\times \sim [A\mathbf{v}]_\times A \quad (7)$$

we have:

$$K_0^T [e]_\times K_0 \sim [K_0^{-1} e]_\times \underbrace{K_0^{-1} K_0}_I \quad (8)$$

Thus, (6) and (8) enable us to write (4) as:

$$\underline{E} \underline{E}^T \sim [K_0^{-1} e]_\times C_\alpha [K_0^{-1} e]_\times^T \quad (9)$$

Let us use the Singular Value Decomposition of \underline{E} :

$$\underline{E} = U \text{diag}(r, s, 0) V^T \quad (10)$$

Introducing (10) into (9), and moving U and U^T to the opposite side of the formula, result in:

$$\text{diag}(r, s, 0) \underbrace{V^T V}_I \text{diag}(r, s, 0) \sim U^T [K_0^{-1} e]_\times C_\alpha [K_0^{-1} e]_\times U \quad (11)$$

(remind a property: $[\mathbf{v}]_\times^T = -[\mathbf{v}]_\times$, for any vector \mathbf{v}). Using (7), we can write (11) in the form:

$$\text{diag}(r^2, s^2, 0) \sim [U^T K_0^{-1} e]_\times U^T C_\alpha U [U^T K_0^{-1} e]_\times \quad (12)$$

Let us notice, that

$$[U^T K_0^{-1} e]_\times = \left[\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right]_\times = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

If we denote with $(\mathbf{u}_1 \ \mathbf{u}_2 \ \mathbf{u}_3)$ columns of matrix U , (12) writes as follows:

$$\begin{pmatrix} r^2 & 0 \\ 0 & s^2 \end{pmatrix} \sim \begin{pmatrix} \mathbf{u}_2^T C_\alpha \mathbf{u}_2 & -\mathbf{u}_1^T C_\alpha \mathbf{u}_2 \\ -\mathbf{u}_1^T C_\alpha \mathbf{u}_2 & \mathbf{u}_1^T C_\alpha \mathbf{u}_1 \end{pmatrix} \quad (13)$$

Equalities between ratios of coefficients of the matrices in (13) form Kruppa's equations. However, only the following equality can contribute positively to the solution:

$$\frac{r^2}{s^2} = \frac{\mathbf{u}_2^T C_\alpha \mathbf{u}_2}{\mathbf{u}_1^T C_\alpha \mathbf{u}_1} \quad (14)$$

The other possible equation ($\mathbf{u}_1^T C_\alpha \mathbf{u}_2 = 0$) leads always to a solution $C_\alpha = I$, and thus $\alpha = 1$.

Remembering the form of C_α (5), one can retrieve from (14) the unknown α , by solving a quadratic equation (since the numerator and denominator of (14) are linear expressions in entries of matrix C_α):

$$\alpha = \sqrt{\frac{s^2 u_{32}^2 - r^2 u_{31}^2}{r^2 (u_{11}^2 + u_{21}^2) - s^2 (u_{12}^2 + u_{22}^2)}} \quad (15)$$

where u_{ij} are entries of matrix U , and r, s - the singular values, given in (10).

4.2. Outline of the algorithm

Step 0: Perform off-line calibration of the camera, obtaining $\alpha_{min}, \alpha_{max}, k, (u_0, v_0)$ - thus full calibration matrices: K_{min} and K_{max} for reference images, and a calibration matrix K (associated with any other image) defined up to α .

Then, for a stereo pair (a reference image and an image of an unknown zooming), given the matching:

Step 1: Compute the fundamental matrix F .

Step 2: Move F to a "semi-calibrated" space, obtaining a new matrix \underline{E} - see (6).

Step 3: Apply the SVD on \underline{E} - see (10).

Step 4: Use entries of the matrices obtained in Step 3 to compute the unknown internal parameter α of matrix K - see (15).

4.3. Critical motions

As it has been fully studied in [14], a solution for the unknown magnification parameter is not always uniquely defined. In our case, since we consider to know all intrinsic parameters of one camera, there exist only one family of camera configurations that is critical, which is significantly less than with more general cases of self-calibration.

Let us consider a stereo pair of cameras: C_{ref} (fully calibrated) and C (with an unknown α). The algorithm is




Figure 3. Critical configuration of cameras

singular if the centre of camera C_{ref} lies on the optical axis of camera C (Figure 3). This kind of configuration is connected with a camera movement (starting from the reference position), that consists of any rotation, followed by a translation along the optical axis of the camera. There is no constraint on the orientation of camera C_{ref} .

The case has been analyzed along the lines of [14]. Here, we omit its explanation, due to the lack of space.

In practice, any camera configuration that is close to the critical one, can cause problems in self-calibration, giving rise to inaccurate results.

5. Matching

We are interested in running our self-calibration algorithm on pairs of images of different spatial resolutions (different magnifications). Being aware of problems concerning matching such views, we propose a way to avoid it, by introducing an additional view, that allows to match the two original ones.

5.1. Difficulties

Existing direct techniques for automatic matching of two images taken from different viewpoints and with different resolutions do not give satisfactory results. Since a big area on the zoomed-in image is to be correlated with a small area on the zoomed-out image, accuracy of computed epipolar lines is weak.

Even when dealing with images very similarly zoomed-in, very few algorithms cope with matching them, if the camera movement between the two views is not small. On the other hand, once one decreases the baseline between cameras (so that it would be appropriate for correlation techniques), the scene reconstruction becomes less reliable.




Figure 4. Connections between images: fundamental matrix F_a between I_{ref} and I_a , affine transformation A between I_a and I_{zoom} , fundamental matrix F between I_{ref} and I_{zoom}

5.2. Our method

To avoid a manual specification of corresponding points, we combine two techniques:

- matching two images of *the same resolution*, taken from *different viewpoints*
- matching two images of *different resolutions*, taken from *the same viewpoint*

Hence, we assume being provided with an additional view I_a , of the same resolution as the reference image I_{ref} , but taken from the same camera position as the zoomed-in one I_{zoom} (see Figure 4).

Having performed an automatic matching [2] between I_{ref} and I_a , we compute the fundamental matrix F_a of that stereo pair. Thus for any image points \mathbf{m}_{ref} , \mathbf{m}_a (related to I_{ref} and I_a respectively):

$$\mathbf{m}_a^T F_a \mathbf{m}_{ref} = 0 \quad (16)$$

On employing method [5], we match I_a with I_{zoom} , and estimate an affine transformation A between them (due to the lack of space, we omit a derivation of this property) - for any image points \mathbf{m}_a , \mathbf{m}_{zoom} (related to I_a and I_{zoom} respectively):

$$\mathbf{m}_a = A \mathbf{m}_{zoom} \quad (17)$$

Now, (16) and (17) let us find out the fundamental matrix F of the stereo pair of interest - I_{ref} and I_{zoom} :

$$(A \mathbf{m}_{zoom})^T F_a \mathbf{m}_{ref} = 0$$

$$\mathbf{m}_{zoom}^T \underbrace{A^T F_a}_{F} \mathbf{m}_{ref} = 0$$

"True α "	Self-calibration for α			
	Calibration grid		Arbitrary object	
	min ref.	max ref.	min ref.	max ref.
708	701±1%	763±8%	731±3%	691±2%
847	814±4%	887±5%	881±4%	829±2%
1018	960±6%	1044±3%	1040±2%	1012±1%
1250	1171±6%	1273±2%	1298±4%	1242±1%
1486	1387±7%	1508±1%	1592±7%	1442±3%
1729	1621±6%	1751±1%	1886±9%	1779±3%
1905	1772±7%	1933±1%	2072±9%	1862±2%

Table 4. Self-calibration results. The same α (in a row, related to a zoom setting) is estimated from stereo pairs of two different objects, with both kind of reference images each. "True α " comes from calibration.

Thus

$$F = A^T F_a \quad (18)$$

Equation (18) enables us to compute the epipolar geometry between images of different resolutions (I_{ref} and I_{zoom}), without being given matches between them. It is sufficient to specify correspondences between each of those images and a special additional one (I_a), and result in connections written as functions of F_a and A .

6. Experiments

Input data. We took images of an arbitrary object (a toy house), with the minimal and the maximal zoom settings, from a reference viewpoint. Then, from another camera position, we shot a number of images, of variable zooming. We also took images of a calibration grid, every time a view of the house was registered. Hence, we had two sets with corresponding images (of different features), taken with identical camera settings (Figure 5).

Separately for each photographed object, we combined our images in stereo pairs of a reference image and an image taken with an unknown zoom. Each image of an unknown zoom was put into 2 stereo pairs: with a minimal and with a maximal zoom reference image. Having employed results from off-line calibration (the constraint: distance camera-scene $\geq 2.5m$ had been fulfilled), we ran the self-calibration algorithm for each stereo pair, obtaining estimations for α , related to every considered zooming (see Table 4).

Discussion. The algorithm recovers the unknown magnification parameter with a high accuracy. However, there are some cases, where the relative error grows up to 9%. They show up for stereo pairs, which combine images of significantly different resolutions (e.g. a minimal zoom reference view with a strongly zoomed-in one; and vice-versa).




Figure 5. Images of different features, taken with identical camera settings.

The reason could be related to the fact, that the considered self-calibration step does not take into account any distortion model, and thus, its results are not always consistent with the off-line calibration (see Section 3.2). In particular: a distortion model, considered for the minimal zoom reference image, is "forwarded" by self-calibration to α , estimated for the other image of the stereo pair. If that image has been taken with a relatively big zoom setting, the no-distortion model has to be considered then, in order to avoid over-parameterization. For the opposite case: not taking distortion into account for the maximal zoom reference case, implies the no-distortion model for the other image, as well, which is not always correct (zoomed-out images).

A way to cope with the described inconsistency would be to employ a non-linear optimization. The self-calibration step, along with a linear structure from motion method, provides an initial guess for camera parameters (internal and external ones). Then, it would be sufficient to use an extended projection model (including distortion) in a bundle-adjustment setting.

Overall, the experiments validate that our self-calibration method is reliable, for any stereo pair. The unknown magnification parameter can be recovered with even 2%-accuracy, provided that the stereo pair is composed of images of similar resolutions. Therefore, it is more convenient to use a minimal zoom reference view to self-calibrate zoomed-out images, and a maximal zoom reference one, for more zoomed-in images.

7. 3D reconstruction

We applied the described technique on a stereo pair of images of a chimney (Figure 6). The only knowledge we had, was that both images were taken with our camera, and that one of them was taken with the minimal possible zoom. Self-calibration provided us with an estimation for the unknown magnification parameter for the second setting of the camera: $\alpha = 980$.

Reliability of the obtained reconstruction of the chim-




Figure 6. Stereo pair of a chimney (the reference image on the left).



Figure 7. Reconstructed chimney.

ney (Figure 7), with correctly retrieved depth and angles between specified planes (Table 5), certifies a high quality of the performed calibration, and thus, capability to recover the Euclidean structure.

8. Conclusion

We have presented a method to simplify the self-calibration process of a zooming camera, based only on information of a boundary (minimal or maximal) zoom, used for taking one of the images. Due to the off-line calibration preprocessing, the on-line self-calibration step has only one parameter to estimate, and thus, only one family of critical motion sequences for cameras to deal with (a situation that is not valid for more complex cases of self-calibration). We provide a close-form solution for the problem and present experiments on real images that validate the stability and reliability of our method.

The proposed combined calibration technique can be easily used in various applications, as quite often one is provided with at least one reference image. The complex problem of dealing with wide, differently zoomed views of a scene, is decomposed into several simpler tasks, which is an important advantage of the presented approach.

Acknowledgements. We would like to thank Frederik Schaffalitzky from the Visual Geometry Group in Oxford, for making accessible a matching software.

Plane	3	5	6	7	8
1	88 (90)	88 (90)	2 (0)	89 (90)	89 (90)
3	-	2 (0)	91 (90)	71 (70)	114 (115)
5	-	-	90 (90)	70 (70)	113 (115)
6	-	-	-	88 (90)	90 (90)
7	-	-	-	-	43 (45)

Table 5. Angles (in [deg]) between chosen planes of the chimney: retrieved values, and the real ones (in parentheses).

References

- [1] M. Armstrong, A. Zisserman, and R. Hartley, "Self-calibration from image triplets", *ECCV*, pp.3-16, 1996.
- [2] A. Baumberg, "Reliable feature matching across widely separated views", *CVPR*, pp. 774-781, 2000.
- [3] S. Bougnoux, "From projective to euclidean space under any practical situation, a criticism of self-calibration", *ICCV*, pp. 790-796, 1998.
- [4] E. Boyer, *Reconstruction de surfaces d'objets courbes en vision par ordinateur*, PhD thesis, INP de Lorraine, 1996.
- [5] Y. Dufournaud, C. Schmid, and R. Horaud, "Matching images with different resolutions", *CVPR*, 2000.
- [6] O. Faugeras, "What can be seen in three dimensions with an uncalibrated stereo rig?", *ECCV*, pp. 563-578, 1992.
- [7] R.I. Hartley, "Estimation of relative camera positions for uncalibrated cameras", *ECCV*, pp. 579-587, 1992.
- [8] R.I. Hartley, "Self-calibration from multiple views with a rotating camera", *ECCV*, pp. 471-478, 1994.
- [9] R.I. Hartley, "Kruppa's equations derived from the fundamental matrix", *PAMI*, 19(2), pp. 133-135, 1997.
- [10] R.I. Hartley, R. Gupta, and T. Chang, "Stereo from uncalibrated cameras", *CVPR*, pp. 761-764, 1992.
- [11] A. Heyden and K. Åström, "Minimal conditions on intrinsic parameters for euclidean reconstruction", *ACCV*, vol. II, pp. 169-176, 1998.
- [12] E. Malis and R. Cipolla, "Multi-view constraints between collineations: Application to self-calibration from unknown planar structures", *ECCV*, pp. 610-624, 2000.
- [13] M. Pollefeys, R. Koch, and L. Van Gool, "Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters", *ICCV*, pp. 90-95, 1998.
- [14] P. Sturm, "Critical motion sequences for the self-calibration of cameras and stereo systems with variable focal length", *BMVC*, pp. 63-72, 1999.
- [15] B. Triggs, "Autocalibration from planar scenes", *ECCV*, 1998.
- [16] R.G. Willson and S.A. Shafer, "Modeling and calibration of zoom lenses", In *Camera Calibration and Orientation Determination*, Springer-Verlag, 1993.
- [17] C. Zeller and O. Faugeras, "Camera self-calibration from video sequences: the Kruppa equations revisited", Technical report 2793, INRIA, 1996.