Online auto-calibration in man-made worlds

Roman Pflugfelder
Video and Safety Technology
ARC Seibersdorf research GmbH
A-2444 Seibersdorf, Austria
roman.pflugfelder@arcs.ac.at

Horst Bischof
Institute for Computer Graphics and Vision
Graz University of Technology
A-8010 Graz, Austria
bischof@icg.tu-graz.ac.at

Abstract

This paper proposes a method to auto-calibrate a static camera in man-made worlds. The calibration is done with vanishing points which are estimated with line segments. Instead of a single image, the method uses constantly acquired images from a camera. This online behavior of the method has two advantages: A continuous and adaptive estimation of the vanishing points over thousands of images suppresses the amount of noise in their positions. Newly detected line segments can further improve the estimated positions of vanishing points. This method is of practical importance, because simple installation and maintenance of cameras in video surveillance is a competitive advantage. Furthermore, applications like tracking can profit, because knowledge about the scene geometry can be a potential clue. The method was evaluated in a broad range of indoor and outdoor environments. The relative error between the internal camera parameter estimates and the true values was on average 5%.

1. Introduction

A rough knowledge about the scene geometry and the camera is a prerequisite for many applications in the area of visual surveillance. For example, length and area ratios are useful to recognize objects. Some object classification methods need image rectification to remove projective and affine distortion. The camera projection matrix allows to project objects in a certain world point onto the image. This is very useful for model-based applications. Some object tracking applications use homographies to operate on a world ground plane. Homographies between cameras or between cameras and a common world ground plane are a powerful cue to track objects in a multi-camera system. Such systems are attractive, because object occlusions can be handled more effectively than in the single camera case.

Unfortunately, to know these homographies or to know

Figure 1. This scene is from the PETS 2001 dataset 3, camera 1. This paper introduces an online EM method that estimates the vanishing points and internal camera parameters simultaneously from line segments (top). The relative error of the focal length to the ground truth was 5.3% after 100 steps. The bottom image shows the orientation of the world coordinate system (black) with respect to the camera coordinate system (white).

Proceedings of the Digital Imaging Computing: Techniques and Applications (DICTA 2005)
0-7695-2467-2/05 $20.00 © 2005 IEEE

IEEE COMPUTER SOCIETY

the camera projection matrix needs information about the scene geometry and the camera. Usually, this information is encoded as prior knowledge in many vision algorithms. This paper concentrates on how to auto-calibrate a single, static camera from vanishing points in man-made environments. Calibration determines the internal parameters of a camera. The vanishing points are simultaneously estimated from line segments, which are aligned in orthogonal world directions.

Auto-calibration is of practical importance, because simple installation and maintenance of cameras in video surveillance is a competitive advantage. State-of-the-art methods for manual calibration use known world points and the corresponding image points. But manual calibration is awkward and time-consuming and needs experience by the user. In the worst case, manual calibration must be repeated each time a change in the environment happens, which makes maintenance rather expensive.

Some other methods use trajectories of moving objects to estimate constraints on the camera projection matrix (Bose [2], Lv [13]). The homography between a world ground plane, e.g. road, and the image plane is estimated using vanishing points and a constant object velocity assumption. Unfortunately, these methods may suffer from inaccurately measured trajectories and inaccurately measured point and line features respectively.

People have also used line segments in single images, which are images of orthogonal lines in the world (Caprile and Torre [4], Rother [14], Kosecka [11]). These line segments intersect in vanishing points which constrain the camera projection matrix. Despite these methods are feasible for video surveillance applications, an adaptation to a changing environment is not considered.

Coughlan [5] proposed to use the image gradients in a Bayesian framework instead of an edge detector. Deutscher [6] showed an interesting calibration approach within this framework. Schindler [15] significantly extended this framework to group line segments and to compute vanishing points. Unfortunately, it is not clear how to automatically initialize their methods.

In contrast to these works, our approach is able to improve the calibration and the estimation of the position of vanishing points if new line segments are present. For example, an object can suddenly appear in the world scene or the illumination conditions improve and permit a more accurate detection of line segments. The adaptive behavior of the approach averages the estimated positions of the vanishing points which suppresses noise. Another aspect of our approach is a plausibility test between the estimated focal length and the focal length of the lens.

Section 2 presents important requirements about the world scene and the camera. Section 3 and section 4 discuss the approach. Experiments and results (section 5) show the applicability of the approach. Section 6 concludes this paper.

2. Requirements of the approach

The approach requires restrictive, environmental characteristics:

1. The mounting of the camera must be rigid and the lens must have a fix focal length.

2. The world scene viewed by the camera must contain orthogonal directions. These directions are given by straight-lined objects that are partly imaged as line segments. Usually, man-made structures are build in an orthogonal manner. In Coughlan [5] these world scenes are called Manhattan worlds.

To handle the Manhattan world assumption with more care, knowledge of the lens and the camera and assumptions about the camera are required:

1. We assume that the image sensor's pixel size and the focal length of the lens is known. In general, the true focal length and the focal length of the lens are not equal. However, a tolerance interval for the true focal length can be given. In our experiments this tolerance was set to 10%.

2. The skew of the image sensor's pixel is assumed to be zero and the aspect ratio of the pixel's side lengths is assumed to be constant and known.

3. In some scenes vanishing points lie close to infinity or only two prominent directions are present. Under such conditions the principal point must be constrained to be close to the image center.

3. Initialization method

The basic information to estimate vanishing points are line segments which are detected from an acquired sequence of images. Our approach is closely related to the work of Kosecka [11] and Rother [14]. Figure 2(a) shows a flow chart of the proposed initialization method.

Especially lenses with wide viewing angles are substantially distorted. Therefore, the acquired image should be rectified using a lens distortion model. The type of the distortion model and the parameters of the model are either known, i.e. a model is chosen and the parameters are precomputed for a specific lens, or the choice of a model and the parameter estimation could be done by an automatic method as suggested by Devernay [7]. Tests have shown that this method works with moderate distortions. It is no solution for heavily distorted images. Practically, to crop

IEEE COMPUTER SOCIETY

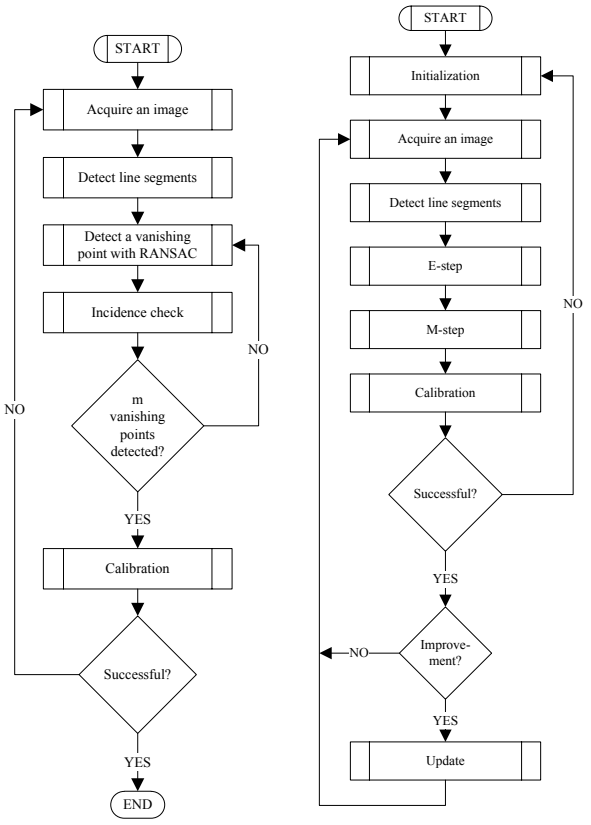

(a) Initialization (b) Refinement

Figure 2. Flow charts of the initialisation and the refinement methods.

the images uniformly at all edges reduces the distortion, because distortion decreases towards the image center.

The Canny detector computes edges in the rectified image. Line segments $\{l_1, \ldots, l_L\}$ are detected as linear edges using a method suggested by Guru [8]. Each line segment $l_j = a_j \times b_j$ is defined by the line segment end points a_j and b_j which are vectors in the homogeneous image coordinates.

The gradient computation in the edge detector is not robust to noise. A preceding averaging over some images substantially improves the accuracy of the gradients.

3.1. Detecting vanishing points

A vanishing point v is an intersection point of a subset of $\{l_1, \ldots, l_L\}$. In imaged Manhattan worlds most line segments have to intersect in a number of vanishing points which encode the dominant and orthogonal directions. Usually, two or three orthogonal directions are present in a world scene ($m \in \{2, 3\}$), but $m > 3$ is also possible.

Kosecka [11] used the orientation of the line segments to detect vanishing points. Peaks of the orientation histogram group $\{l_1, \ldots, l_L\}$ into m sub sets. Unfortunately, line segments with the same orientation can also intersect in different vanishing points in some situations.

Our method avoids this problem by using RANSAC (Random Sample Consensus) and is similar to the work of Rother [14]. The simple idea of RANSAC is to construct repeatedly a vanishing point $v = l_i \times l_j$. l_i and l_j are chosen randomly from $\{l_1, \ldots, l_L\}$. A line segment $l_k \in \{l_1, \ldots, l_L\} - \{l_i, l_j\}$ will be an inlier, if l_k meets v within a given error σ. RANSAC tries to find a vanishing point v where a maximal number of line segments are inliers.

Generally, line segments are uncertain due to noisy images. Hence, a first order error analysis is done in all estimations. See Heuel [10] for a rigorous discussion among this topic. Following Liebowitz [12], the isotropic noise in the end points a_j and b_j is modeled as Gaussian random variable ξ

$$a_j = \bar{a_j} + \xi, b_j = \bar{b_j} + \xi, \xi \sim \mathcal{N}(0, \sigma I_{3\times 3}), \qquad (1)$$

where $\bar{a_j}$ and $\bar{b_j}$ are the true end points and I is the 3×3 identity matrix.

The error of l_j meeting v can be formulated with this simple noise model. In general, l_j will never meet v perfectly due to noise. However, a straight line $\bar{l_j}$ can be constructed which will meet v exactly. See figure 3 for the geometric details. Liebowitz showed that $\bar{l_j}$ is the MLE of the true line segment $\bar{a_j} \times \bar{b_j}$. The errors for a_j and b_j are the distances $d(a_j, \bar{a_j})$ and $d(b_j, \bar{b_j})$.

In contrast to Rother, RANSAC is used consecutively to detect the vanishing points. To avoid that a vanishing point v_i lies within the uncertainty of a previously detected vanishing point v_j a test statistics T is computed by

$$T = (v_i - v_j)^\top \Sigma_j^{-1} (v_i - v_j) \qquad (2)$$

where Σ_j are the variance co-variance matrices of v_j. If no incidence happens, the set of line segments will be reduced by the line segment inliers and RANSAC is repeated. Experiments have further shown that RANSAC will perform more robust, if line segments are chosen according to their pixel length.

3.2. Calibration

Vanishing points $v_1, \ldots, v_m$ give $\binom{m}{2}$ constraints on the IAC $\omega_{3\times 3}$ (Image of the Absolute Conic). ω can be interpreted as a metric of the uncalibrated space. ω is a symmetric matrix and has 5 degrees of freedom. The textbook of Hartley and Zisserman [9] provides more information about the IAC. Furthermore, our assumptions of zero skew s and a

IEEE COMPUTER SOCIETY

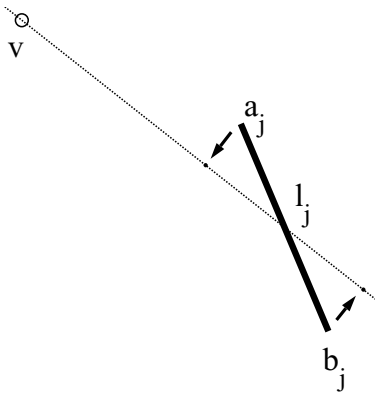

Figure 3. Geometric illustration: The line segment l_j does not meet exactly the vanishing point v.

constant aspect ratio r yield two more constraints on ω. To assume that the principal point $p = (u_0\ v_0)^\top$ is close to the image center gives further two constraints on ω. All these constraints can be formulated with the following equations:

$$\begin{aligned} v_i^\top \omega v_j &= 0 \quad 1 \geq i, j \geq m,\ i \neq j & (3)\\ (1\,0\,0)^\top \omega (0\,1\,0) &= 0 & (4)\\ (1\ r\ 0)^\top \omega (0\ -1\ 0) &= 0 & (5)\\ \omega p &= (0\,0\,1)^\top & (6) \end{aligned}$$

For $m > 1$ a solution of ω can be computed. If more than two vanishing points are detected, the linear equations system formed by the equations in 3-6 is over-determined. An optimal solution of ω in a least-squares sense can then be computed using SVD (Singular Value Decomposition).

To know ω is equal to know the internal camera parameters

$$\omega = K^{-\top} K^{-1}, \tag{7}$$

with

$$K = \begin{pmatrix} f & s & u_0 \\ 0 & rf & v_0 \\ 0 & 0 & 1 \end{pmatrix}.$$

f is the focal length. K can be computed using a Cholesky decomposition of ω.

A solution for K can be incorrect, because information about the orthogonality of world directions is lost during the imaging process. A simple way to test the plausibility of a solution of K is to test the relative error between f and the focal length of the lens f_0. If the relative error is smaller than a maximal error threshold ϵ

$$\frac{f - f_0}{f_0} < \epsilon, \tag{8}$$

the solution of K is accepted, otherwise it is rejected and the initialization method is repeated. This test is sufficient for a correct solution of K, if the principal point is constrained. The condition that ω is positive definite and that a Cholesky decomposition is possible is necessary but not sufficient.

4. Refinement method

Figure 2(b) shows a flow chart of the proposed refinement method. The problem now is to group line segments in new images that meet the current vanishing point estimates. Simultaneously, these vanishing point estimates should be re-estimated using the expected grouping. This very general problem can be solved with the EM algorithm (Expectation Maximization). The RANSAC-based initialization method (see section 3) gives an initial guess of the vanishing point estimates. The refinement method uses an online version of the EM algorithm which is discussed in Brochu [3].

4.1. Likelihood function

The likelihood of each l_i to meet a v_k can be written as a weighted mixture of likelihood functions

$$\sum_{k=1}^{m} \Pr(v_k) p(l_i | v_k, \sigma) + \Pr(noise) p(noise|\sigma). \tag{9}$$

$\Pr(v_k)$ are the prior probabilities of v_k. These priors are initialized with the relative frequency between line segments meeting v_k and the total number L of line segments detected. If l_i does not meet any v_k, it will be a noisy line segment. The prior probability of the occurrence of noisy line segments is always

$$\Pr(noise) = 1 - \sum_{k=1}^{m} \Pr(v_k). \tag{10}$$

We assumed, that the likelihood value $p(noise|\sigma)$ can be evaluated with a one-sided Gaussian pdf

$$p(noise|\sigma) = \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(-\frac{1}{\sigma}\right) \tag{11}$$

at location 2σ, i.e. line segments meeting non of the v_k within 2σ are more likely to be noisy. The likelihood $p(l_i|v_k, \sigma)$ can be expanded to

$$p(l_i|v_k, \sigma) = p(a_i|v_k, \sigma) p(b_i|v_k, \sigma), \tag{12}$$

because a_i and b_i are independent from each other. $p(a_i|v_k, \sigma)$ and $p(b_i|v_k, \sigma)$ are one-sided Gaussian pdfs

$$\begin{aligned} p(a_i|v_k, \sigma) &= \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(\frac{-d^2(a_i, \bar{a_i})}{2\sigma^2}\right), & (13)\\ p(b_i|v_k, \sigma) &= \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(\frac{-d^2(b_i, \bar{b_i})}{2\sigma^2}\right), & (14) \end{aligned}$$

IEEE COMPUTER SOCIETY

where $d^2(a_i, \bar{a_i})$ and $d^2(b_i, \bar{b_i})$ are the squared distances between the end points a_i, $\bar{a_i}$ and b_i, $\bar{b_i}$ respectively. Now, the online EM algorithm tries to find optimal estimates of $v_1, \ldots, v_m$ and $\Pr(v_1), \ldots, \Pr(v_m)$ under the given l_i and a given σ.

4.2. E-step

The goal of the E-step is to compute the posterior membership probability

$$q_{ki} = \Pr(v_k|l_i) = \frac{p(l_i|v_k, \sigma)\Pr(v_k)}{p(l_i)}, \tag{15}$$

that l_i meets v_k, given the prior probabilities $\Pr(v_k)$ and the likelihood function $p(l_i|v_k, \sigma)$. $p(l_i)$ is a normalization factor to ensure

$$\sum_{k=1}^{m} q_{ki} + \Pr(l_i|noise) = 1, \tag{16}$$

and can be written as

$$p(l_i) = \sum_{k=1}^{m} p(l_i|v_k, \sigma)\Pr(v_k) + p(noise|\sigma)\Pr(noise). \tag{17}$$

$\Pr(l_i|noise)$ is the probability that line segment l_i is noisy and can be computed from equation 16. The E-step gives us the best guess of the membership q_{ki} of unknown line segments l_i to the current vanishing point estimates v_k with prior probabilities $\Pr(v_k)$. If some l_i do not meet any of the v_k, they are expected to be noisy line segments.

4.3. M-step

In the M-step the current vanishing point estimates v_k are re-estimated using the previously computed membership probabilities. This is done by maximizing the expected log-likelihood function

$$\max_{v_k^*} J(v_k^*) = q_{ki} \log p(l_i|v_k^*, \sigma) \tag{18}$$

with respect to new estimates of vanishing points v_k^*. Section 4.1 showed that $p(l_i|v_k^*, \sigma)$ and consequently $\log p(l_i|v_k^*, \sigma)$ are likelihood functions based on geometric error distances between l_i and its MLE with respect to v_k^*. Unfortunately, no explicit and optimal estimate of v_k^* can be given, because $J(v_k^*)$ is a non-linear function. However, iterative, numerical algorithms like Levenberg-Marquardt can be used to find optimal estimates of v_k^*. We refer to Liebowitz [12, 3.6, p.63-69], who shows in all details the MLE of line segment intersections.

Now, the old estimate v_k^{t-1} is adapted in terms of weighted-means with the new estimates v_k^* by

$$v_k^t = (1 - \lambda)v_k^{t-1} + \lambda v_k^*, \tag{19}$$

where λ is a constant learning rate.

Similarly, the prior probabilities are adapted by

$$\Pr(v_k)^t = (1 - \lambda)\Pr(v_k)^{t-1} + \lambda \frac{\sum_{i=1}^{L} q_{ki}}{L} \tag{20}$$

with the number of line segments L.

4.4. Update and calibration

Line segments can disappear or line segments can be detected during the adaptation, because of changes in illumination or changes in the structure of the scene. Therefore, it is a good idea to keep the best vanishing point estimates so far or to update worse estimates. The same is true for the internal camera parameters which are estimated from the vanishing points. Let $\hat{v}_k$ be the vanishing points which should be used for online calibration. In the beginning, each $\hat{v}_k$ is set to v_k^0. We suggest to use the uncertainty in the entries of $\hat{v}_k$ and v_k^t to decide about an update. As mentioned in section 3.1, a first order error analysis provides variance co-variance matrices $\Delta_{\hat{v}_k}$ and $\Delta_{v_k^t}$. The elements of $\Delta_{v_k^t}$ will be smaller, if more and longer line segments are used to compute v_k^t. Contrary, if the angle between two line segments intersecting in a vanishing point is small, the uncertainty of the vanishing point will be large. Otherwise, if the angle between two line segments is large ($< \frac{\pi}{2}$) the uncertainty of their intersection point is relatively smaller. Look at Liebowitz [12, 3.6, p.72-77] for a detailed discussion of the uncertainty of intersection points.

As a measurement of the uncertainty of $\hat{v}_k$ and v_k^t we use the trace of their variance co-variance matrices. If $trace(\Delta_{v_k^t}) < trace(\Delta_{\hat{v}_k})$, $\hat{v}_k$ will be updated with v_k^t. Otherwise, no update happens.

The calibration can fail during the EM iterations. If vanishing points lie far apart of the image, the Cholesky decomposition may fail, because of the bad conditioning of ω. This degenerate case can be avoided by slightly rotating the camera. If the camera is suddenly rotated, the current vanishing point estimates will no longer be valid. This situation can easily happen during a maintenance of the camera, where e.g. the camera is cleaned. If slightly wrong line segments appear in the image, the online EM algorithm will update the vanishing points with invalid optimal estimates. For example, a new object appears in the world scene and its alignment to the current directions varies slightly to the directions of objects seen so far. In all cases, the initialization method should be re-started.

5. Experiments and results

In all experiments we set σ to 0.25pixel, and ϵ to 0.1. λ was set to 0.01. Each line segment was at least 15pixels long.

IEEE COMPUTER SOCIETY

First we tested the initialization method in a broad range of man-made world scenes and it worked successfully in world scenes like public places, corridors and buildings. For example, figure 4(a) shows a public place next to a metro station. The image was taken with a Lumenera 125C and a Tamron lens with a focal length of 3731.3 pixel. The camera was successfully calibrated with three detected vanishing points in the first iteration. To prove these results qualitatively, an overhead view of the ground plane was generated with the vanishing point estimates. The origin of the world coordinate system was imaged to the principal point estimate. Figure 4(b) shows a satisfying metric rectification of the ground plane.

Another successful initialization is shown in figure 5(a) and 5(b). The scene is taken from the Tau-dance data set and shows view 1. Figure 1 shows a successful initialization and refinement on images of the PETS 2001 data set.

In the next experiment we demonstrate the advantages of the online behavior. Figure 6(a) shows our office taken with a Logitech QuickCam Pro4000. The ceiling light in the right upper corner of the image is on. The light hides important edge information. Figure 6(b) shows an image, 300 image frames later. The light is now off, also the door is shut. The edge information is now apparent. Figure 5 shows the adaptation of the internal camera parameters. It can be seen that new apparent line segments can improve the calibration of the internal camera parameters. Table 1 shows the internal camera parameter estimates in detail.

6. Conclusion

We showed how to auto-calibrate a camera in man-made environments. The online behavior of the initialization method and the successive refinement method showed advantages in contrast to existing approaches. To take new line segment information into account can significantly improve the estimation of the internal camera parameters. To test the focal length estimate with the focal length of the lens makes that the estimates are plausible. One improvement for the future could be to use gradient information directly instead of line segments. Our future work will concentrate on using the vanishing points and the calibration to compute the fundamental matrix between two slightly overlapping fields of view.

7. Acknowledgement

This work is part of the Forschungsförderungsfond (FFF) project, *Plug&Detect: Self-calibrating and self-configuring video surveillance*. We appreciate to use the Tau-dance data set from our MUSCLE[1] partner [1]. We

[1] http://www.muscle-noe.org/

(a) Public place

(b) Overhead view

Figure 4. Qualitative analysis of the initialization. (a) shows a public place close to a metro station. All detected line segments are drawn in white color. (b) shows the generated overhead view.

IEEE COMPUTER SOCIETY

(a) Tau-dance scene, view 1

(b) Overhead view

Figure 5. Qualitative analysis of the initialization. (a) shows the view of camera 1. All detected line segments are drawn in white color. (b) shows the generated overhead view within the black rectangle drawn in (a).

(a) Light on

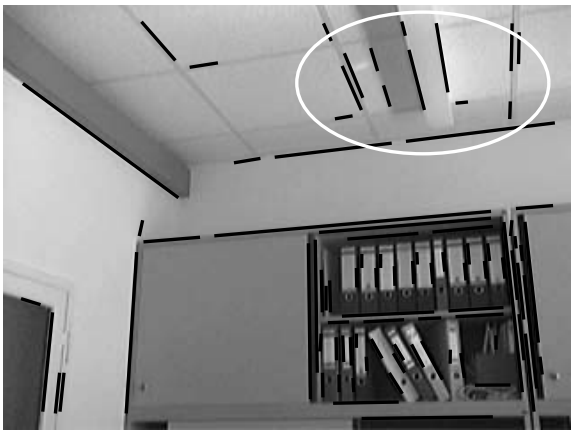

(b) Light off

Figure 6. Illumination experiment at the office. (a) shows an image frame under bad illumination conditions. Consider the upper, right corner where much illumination is present caused by the ceiling light. All detected line segments are drawn in white color. (b) shows the same scene after acquiring 300 new image frames. The lamp was turned off and many more line segments are suddenly present. Note that the door was also shut.

IEEE COMPUTER SOCIETY

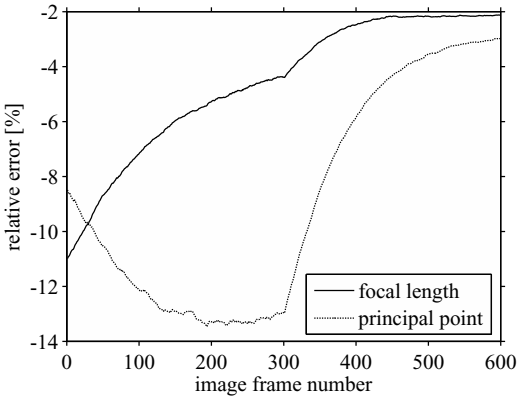

Figure 7. The improvement of the focal length and the principal point over 600 iterations is shown. Note the sharp bend at image frame number 300.

1	[pixel]	variances and co-variances [pixel]				
f	355.6	9079.4	0.0	9079.4	-2311.2	-5954.4
k	-0.0	0.0	0.0	0.0	0.0	-0.0
rf	355.6	9079.4	0.0	9079.4	-2311.2	-5954.4
u_0	132.7	-2311.2	0.0	-2311.2	2427.5	1639.1
v_0	129.1	-5954.4	-0.0	-5954.4	1639.1	4111.8
300	[pixel]	variances and co-variances [pixel]				
f	382.0	128.3	0.0	128.3	-45.4	-90.1
k	0.0	0.0	0.0	0.0	-0.0	-0.0
rf	382.0	128.3	0.0	128.3	-45.4	-90.1
u_0	127.6	-45.4	0.0	-45.4	51.4	33.1
v_0	121.3	-90.1	-0.0	33.1	33.1	70.7
600	[pixel]	variances and co-variances [pixel]				
f	390.0	94.4	0.0	94.4	-31.1	-65.5
k	0.0	0.0	0.0	0.0	-0.0	-0.0
rf	390.0	94.4	0.0	94.4	-31.1	-65.5
u_0	154.9	-31.1	0.0	-31.1	25.0	21.4
v_0	118.2	-65.5	-0.0	-65.5	21.4	50.2

Table 1. Camera calibration results of the office scene. The ground truth for comparison: $f = 399.5$, $k = 0$, $rf = 399.8$, $u_0 = 162.3$, $v_0 = 120.7$. The second column shows the internal camera parameter estimates after 1, 300 and 600 iterations of the calibration. Columns 3-7 show the variances and co-variances of the estimates in column 1. Note the improvement of these uncertainties caused by the averaging in the refinement. After image frame 300 the improvement is also caused by the presence of new line segments.

would like to thank Michael Nölle, Gustavo Fernandéz and Matthias Függer for their helpful discussions.

References

[1] L. Bar, S. Rochel, and N. Kiryati. Tau-dance: Tel-aviv university multiview and omnidirectional video dance database. VIA - Vision and Image Analysis Laboratory, School of Electrical Engineering, Tel Aviv University, January 2005.

[2] B. Bose and E. Grimson. Ground plane rectification by tracking moving objects. In *Proceedings of the Joint IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance (VS-PETS)*. IEEE Computer Society, October 2003.

[3] E. Brochu, N. de Freitas, and K. Bao. Owed to a martingale: A fast bayesian on-line em algorithm for multinomal models. Technical report, University of British Columbia, 2004.

[4] B. Caprile and V. Torre. Using vanishing points for camera calibration. *International Journal of Computer Vision*, 4:127–140, 1990.

[5] J. M. Coughlan and A. Yuille. Manhattan world: Compass direction from a single image by bayesian inference. In *Proceedings of the International Conference on Computer Vision*, volume 2, page 941. IEEE Computer Society, 1999.

[6] J. Deutscher, M. Isard, and J. MacCormick. Automatic camera calibration from a single manhattan image. In *Proceedings of the 7th European Conference on Computer Vision*, volume 2353, page 175. IEEE Computer Society, May 2002.

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

[8] D. S. Guru, B. H. Shekar, and P. Nagabhushan. A simple and robust line detection algorithm based on small eigenvalue analysis. *Pattern Recognition Letters*, 25(1):1–13, 2004.

[9] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[10] S. Heuel. *Uncertain Projective Geometry*. Springer Verlag GmbH, 2004.

[11] J. Kosecka and W. Zhang. Video compass. In *Proceedings of the 7th European Conference on Computer Vision*, volume 2353, page 476, May 2002.

[12] D. Liebowitz. *Camera Calibration and Reconstruction of Geometry from Images*. PhD thesis, University of Oxford, 2001.

[13] F. Lv, T. Zhao, and R. Nevatia. Self-calibration of a camera from video of a walking human. In *Proceedings of the International Conference on Computer Vision*. IEEE Computer Society, August 1999.

[14] C. Rother. A new approach for vanishing point detection in architectural environments. In *Proceedings of the British Machine Vision Conference*, volume 20, pages 647–656, 2002.

[15] G. Schindler and F. Dellaert. Atlanta world: An expectation maximization framework for simultaneous low-level edge grouping and camera calibration in complex man-made environments. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 203–209, 2004.

IEEE COMPUTER SOCIETY