Computer Vision Winter Workshop 2006, Ondřej Chum, Vojtěch Franc (eds.)
Telč, Czech Republic, February 6–8
Czech Pattern Recognition Society

Computation of the epipolar geometry in slightly overlapping views

Pflugfelder Roman[1] and Bischof Horst[2]

[1] Video and Safety Technology
ARC Seibersdorf research
A-2444 Seibersdorf, Austria
roman.pflugfelder@arcs.ac.at

[2] 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 compute the epipolar geometry from slightly overlapping views. Instead of using point correspondences alone, we additionally use the infinite homography. Interestingly, the infinite homography can be determined from line segments when we have the same alignment of the scene in both views. Thereby, vanishing points and the intrinsic camera parameters are computed when we have a Manhattan world scene and assume cameras with square pixels. All assumptions are not too restrictive and occur frequently. The infinite homography and one of the epipoles determine the fundamental matrix. The paper demonstrates a robust and automatic detection of vanishing points and a simultaneous calibration. It further presents a simple solution to compute the epipoles from at least two point correspondences without requiring points in general position. Experiments on real images confirm the applicability of the idea.*

1 Introduction

Robust and automatic methods for the computation of the fundamental matrix are known [8][20]. These algorithms use putative point correspondences to compute the fundamental matrix. Such point correspondences are typically detected by matching image features - even in wide-baseline camera settings [1][15]. These methods give accurate results. However, the accuracy depends on the amount of point correspondences in general position, i.e. points should be distributed uniformly. Unfortunately, cameras with only slightly overlapping views provide point correspondences in a small part of their images. The image in Figure 1 does not provide even that, correspondences can only be established by the person in the room. This work confines itself to man-made environments with orthogonal and parallel structures. These Manhattan worlds contain line segments which determine vanishing points of orthogonal directions in each view [4]. Corresponding vanishing points and the square pixel assumption define the infinite homography [8]. The square pixel assumption is very general and holds for the most digital cameras [14]. It is shown in [8] that a homography and at least two point correspondences determine the fundamental matrix.

Figure 1: Two views of an office scene. The computed epipolar geometry is shown by points and corresponding epipolar lines.

This paper shows that an accurate computation of the fundamental matrix in slightly overlapping views is possible. The infinite homography encapsulates the rotation between the two cameras and their intrinsic parameters. Line segments compensate for the lack of point correspondences in general position and allow the recovery of rotation and intrinsic parameters. At least two point correspondences

within the overlapping image area and the infinite homography determine the epipoles. Experiments with real images show that the proposed method gives reasonable, qualitative results, even if only a few point correspondences are known. It is simple and linear and can be efficiently implemented.

The background of this method is its application in visual surveillance, e.g. to handover objects between two views [3]. The handovering will at least be simplified, if the epipolar geometry is known. Until now, several methods to compute a homography induced by the ground-plane exist [10][18]. However, all these methods assume substantial, overlapping views. Only the method of Khan and Shah [11] would also work in slightly overlapping views. Their work use foot points of tracked people to compute the border lines of the fields of two overlapping views. In contrast to this work, the proposed method does not assume that any point correspondences have to lie on the ground-plane. Such point correspondences should in future be generated by a people tracker.

The paper is organized as follows: Section 2 discusses the computation of vanishing points from line segments and the computation of the intrinsic parameters. Section 3 treats the proposed method to compute the fundamental matrix. After showing and discussing the experiments with synthetic images and the images of an office scene (section 4), the paper concludes with section 5.

2 Vanishing points and calibration

The basic information to estimate vanishing points are line segments which are independently detected from an acquired sequence of images in every single camera view. Our approach is closely related to the work of Kosecka [12] and Rother [16].

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 [5]. Tests have shown that this method works with moderate distortions. Practically, to crop 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 [7]. 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.

2.1 Detecting vanishing points

A vanishing point v is an intersection point of a sub-set 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.

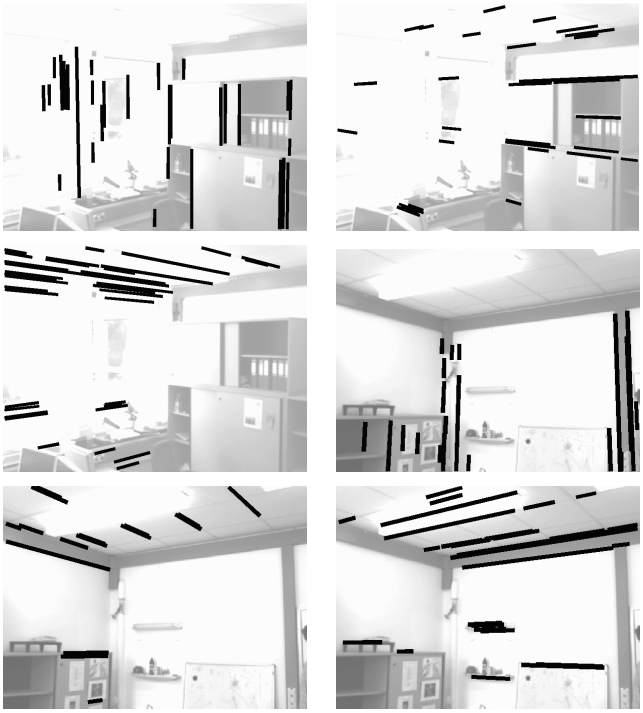

Figure 2: Detected vanishing points in the office scene.

Kosecka [12] 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 [16]. 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 [9] for a rigorous discussion among this topic. Following Liebowitz [13], 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

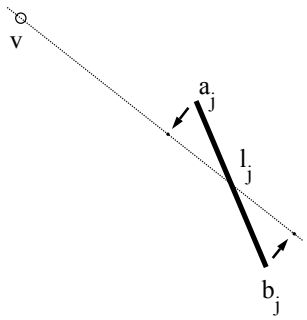

Figure 3: Geometric illustration: The line segment l_j does not meet exactly the vanishing point v.

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

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

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. Figure 2 shows the detected vanishing points in the office scene of Figure 1. Experiments have further shown that RANSAC will perform more robust, if line segments are chosen according to their pixel length.

2.2 Calibration

Vanishing points $v_1, \ldots, v_m$ give $\binom{m}{2}$ constraints on the IAC $\omega_{3\times3}$ (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 [8] provides more information about the IAC. Furthermore, our assumptions of zero skew s and a 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 intrinsic 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.

2.3 Refinement

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 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 [2].

2.4 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

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

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.

2.5 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.

2.6 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 2.4 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 [13, 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.

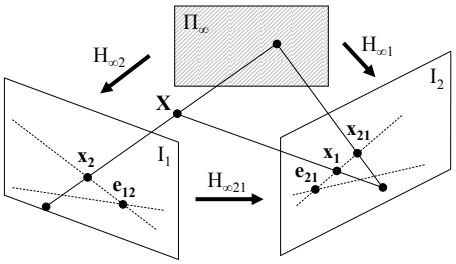

Figure 4: A sketch of the epipolar geometry.

2.7 Correspondence of vanishing points

The correspondences between the detected vanishing points of the same world directions in two views are manually defined. In future this step should also work automatically. Hartley [8] and Rother [17] then show the computation of the camera's rotation - in the remainder of the paper R_1 and R_2 - from vanishing points and the intrinsic parameters of both cameras - in the remainder of the paper K_1 and K_2.

3 Proposed method

The fundamental matrix F completely encapsulates the epipolar geometry between the images of two views. See the illustration in Figure 4. It is known that F can be computed from at least seven point correspondences in these images. Corresponding points are popular. Methods like the eight point algorithm does not deliver accurate results. However, another less known idea exists which also works in slightly overlapping views.

Theorem 1. *From [6]: One of the epipoles e_{21} in image I_1 or e_{12} in image I_2 and every plane induced homography such as the infinite homography $H_{\infty 21}$ or its dual $H_{\infty 12} = H_{\infty 21}^{-1}$ determine F. More formally,*

$$F \simeq [e_{21}]_\times H_{\infty 21} \simeq H_{\infty 21}^{-\top}[e_{12}]_\times \simeq H_{\infty 12}^{\top}[e_{12}]_\times. \tag{21}$$

$[\cdot]_\times$ is the skew-symmetric matrix operator defined in [6]. The $\simeq$ comparator denotes similarity up to a scalar factor due to the use of homogeneous coordinates. We will now show, that $H_{\infty 21}$ can be computed from R_1, R_2, K_1 and K_2. e_{21} is then computed with $H_{\infty 21}$ and at least two point correspondences that do not necessarily have to lie in general position. Analogously, the same is valid for $H_{\infty 12}$ and e_{12}.

3.1 Infinite homography

The infinite homography $H_{\infty 21}$ is the homography from image I_2 to image I_1 induced by the plane at infinity Π_∞. $H_{\infty 21}$ maps a vanishing point v_2 in I_2 to a vanishing point v_1 in I_1 with $v_1 \simeq H_{\infty 21}v_2$. Basically, $H_{\infty 21}$ can be decomposed in a homography $H_{\infty 2}^{-1}$ from I_2 to Π_∞, a rotation R and a second homography $H_{\infty 1}$ from Π_∞ to I_1. $H_{\infty 1}$ and $H_{\infty 2}$ map a direction in the scene to a vanishing point in I_1 with $v_1 \simeq H_{\infty 1}d_1$ and I_2 with $v_2 \simeq H_{\infty 2}d_2$ respectively. d_1 and d_2 are the same direction in the scene but are different coordinate vectors, because the cameras are differently aligned in the scene. If we assume an identical alignment of the scene in both views and as the alignment of the cameras

is already known by R_1 and R_2, the rotation between d_1 and d_2 will be $d_1 = Rd_2$ with

$$R = R_1 R_2^{-1}. \tag{22}$$

This assumption is valid in many scenes of a Manhattan world and can be used in many practical applications. Corollary 2 follows immediately from this assumption.

Corollary 2. *If the identical alignment assumption is valid and R_1, R_2, K_1 and K_2 are known, then*

$$H_{\infty 21} \simeq H_{\infty 1} R H_{\infty 2}^{-1} \simeq K_1 R_1 R_2^{-1} K_2^{-1}. \tag{23}$$

3.2 Epipoles

The infinite homography holds the rotational and the intrinsic information of the epipolar geometry. The only missing information is the translation between the views which is given by the epipoles. The epipole e_{21} in I_1 is the non-trivial left null-space solution of $F^\top e_{21} = 0$. Equivalently, e_{12} in I_2 is the non-trivial right null-space solution of $Fe_{12} = 0$. Consequently, e_{21} must lie on every epipolar line Fx_2 for all points x_2 in I_2. Analogously, the same applies for e_{12}. Conversely, at least two known epipolar lines will define the epipole. In [8, 13.3,336] it was shown that a homography and a pair of corresponding points define the corresponding epipolar line. The reason for that is called plane induced parallax as illustrated schematically in Figure 4. Consider without loss of generality the case where this homography is $H_{\infty 21}$. The ray through one of the corresponding points maps to its corresponding epipolar line. Because all points on the ray lie on the epipolar line, the ray's direction which is the intersection of the ray with Π_∞ must also map onto the epipolar line.

Definition 1. *Let $\{(x_1^i, x_2^i) \mid 1 \leq i \leq N, N \geq 2\}$ be tuples of corresponding points. Let all x_2^i map to x_{21}^i in I_1 with $x_{21}^i \simeq H_{\infty 21} x_2^i$. A measurement matrix M can be constructed with*

$$M = \begin{pmatrix} x_1^1 \times x_{21}^1 & \cdots & x_1^N \times x_{21}^N \end{pmatrix}^\top. \tag{24}$$

The result is that M can be used to compute the epipoles.

Lemma 3. *The epipole e_{21} is the non-trivial solution of the right null-space of $Me_{21} = 0$.*

The problem of estimating the intersection of N lines is a least-squares problem and can be solved using Singular Value Decomposition. The problem is shown with proofs in more detail in [6] and [8]. Analogously, the same applies for e_{12}. Instead of $H_{\infty 21}$ we use $H_{\infty 12} = H_{\infty 21}^{-1}$. Alternatively, e_{12} is also given by equation (21) with $[e_{12}]_\times = H_{\infty 21}^\top F$.

4 Experiments

We tested the proposed method on a real office scene. Although no overlap of the background exists, the fields of view of the cameras overlap substantially. Manually defined points and their corresponding epipolar lines are shown in Figure 1. We used the person in the images to get two point correspondences. A fundamental matrix was also computed with the self-calibration method discussed in [19] which uses hundreds of point correspondences produced by a moving light source. To compare fundamental matrices with each other we used the error measure of Zhang [6, p. 338-339]. The difference of our fundamental matrix compared to the result of the self-calibration method was only 6.42pixel.

In a second experiment, we computed the fundamental matrix between two images of a corridor scene (see Figure 5). The scene is difficult, because occlusion and the perspective distortion allows only a few reliable point correspondences. Three point correspondences (green) were used to compute the epipoles. A qualitative test on other point correspondences suggests a reasonable result for the fundamental matrix.

5 Conclusion

This paper presented a method to compute the fundamental matrix in slightly overlapping views. The idea is to compute the infinite homography from vanishing points and the intrinsic parameters and the epipoles from point correspondences not necessarily in general position. The vanishing points and the intrinsic parameters can be simultaneously computed from line segments. The necessary assumptions are not too restrictive and are valid in many real scenes. Experiments on a real office scene and a real corridor scene confirm the applicability of the proposed method and show reasonable results. Our future work will concentrate towards an automatic matching of vanishing points in different views. Moreover, the computation of the epipolar geometry should adapt to a changing environment.

Acknowledgement

This work is part of the Forschungsförderungsfond (FFF) project 809.674, *Plug&Detect: Self-calibrating and self-configuring video surveillance*.

References

[1] Herbert Bay, Vittorio Ferrari, and Luc Van Gool. Wide-baseline stereo matching with line segments. In *Proceedings of the International Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 329–336, San Diego, US, June 2005. IEEE.

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

[3] Ting-Hsun Chang and Shaogang Gong. Tracking multiple people with a multi-camera system. In *IEEE Workshop on Multi-Object Tracking*, 2001.

[4] James M. Coughlan and A.L. Yuille. Manhattan world: Compass direction from a single image by bayesian inference. In *Proceedings of the International Conference on Computer Vision (ICCV)*, volume 2, pages 941–947, Corfu, Greece, September 1999. IEEE.

[5] Frederic Devernay and Olivier Faugeras. Straight lines have to be straight. *Machine Vision and Applications*, 13(1):14–24, 2001.
[6] Olivier Faugeras and Quang-Tuan Luong. *The Geometry of Multiple Images*. MIT Press, 2001.
[7] 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.
[8] Richard Hartley and Anrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.
[9] Stefan Heuel. *Uncertain Projective Geometry*. Springer Verlag GmbH, 2004.
[10] Christopher Jaynes. Multi-view calibration from planar motion trajectories. *Image and Vision Computing*, 22(7):535–550, July 2004.
[11] Sohaib Kahn and Mubarak Shah. Consistent labeling of tracked objects in multiple cameras with overlapping fields of view. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1355–1360, October 2003.
[12] J. Kosecka and W. Zhang. Video compass. In *Proceedings of the 7th European Conference on Computer Vision (ECCV)*, volume 2353, page 476f. Springer-Verlag, May 2002.
[13] David Liebowitz. *Camera Calibration and Reconstruction of Geometry from Images*. PhD thesis, University of Oxford, 2001.
[14] David Liebowitz and Andrew Zisserman. Combining scene and auto-calibration constraints. In *Proceedings of the International Conference on Pattern Recognition (ICCV)*, volume 1, pages 293–300. IEEE, September 1999.
[15] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide baseline stereo from maximally stable extremal regions. In *Proceedings of the British machine vision conference*, volume 1, pages 384–393, London, 2002. Stephens & George Print Group.
[16] Carsten 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.
[17] Carsten Rother. *Multi-View Reconstruction and Camera Recovery using a Real and Virtual Reference Plane*. PhD thesis, Royal Institute of Technology, January 2003.
[18] Chris Stauffer, Kinh Tieu, and Lily Lee. Robust automated planar normalization of tracking data. In *Joint IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance*, pages 259–266. IEEE, 2003.
[19] Tomas Svoboda, Daniel Martinec, and Tomas Pajdla. A convinient multi-camera self-calibration for virtual environments. *PRESENCE: Teleoperators and Virtual Environments*, 14(4):407–422, August 2005.
[20] P.H.S Torr and Andrew Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78:138–156, 2000.

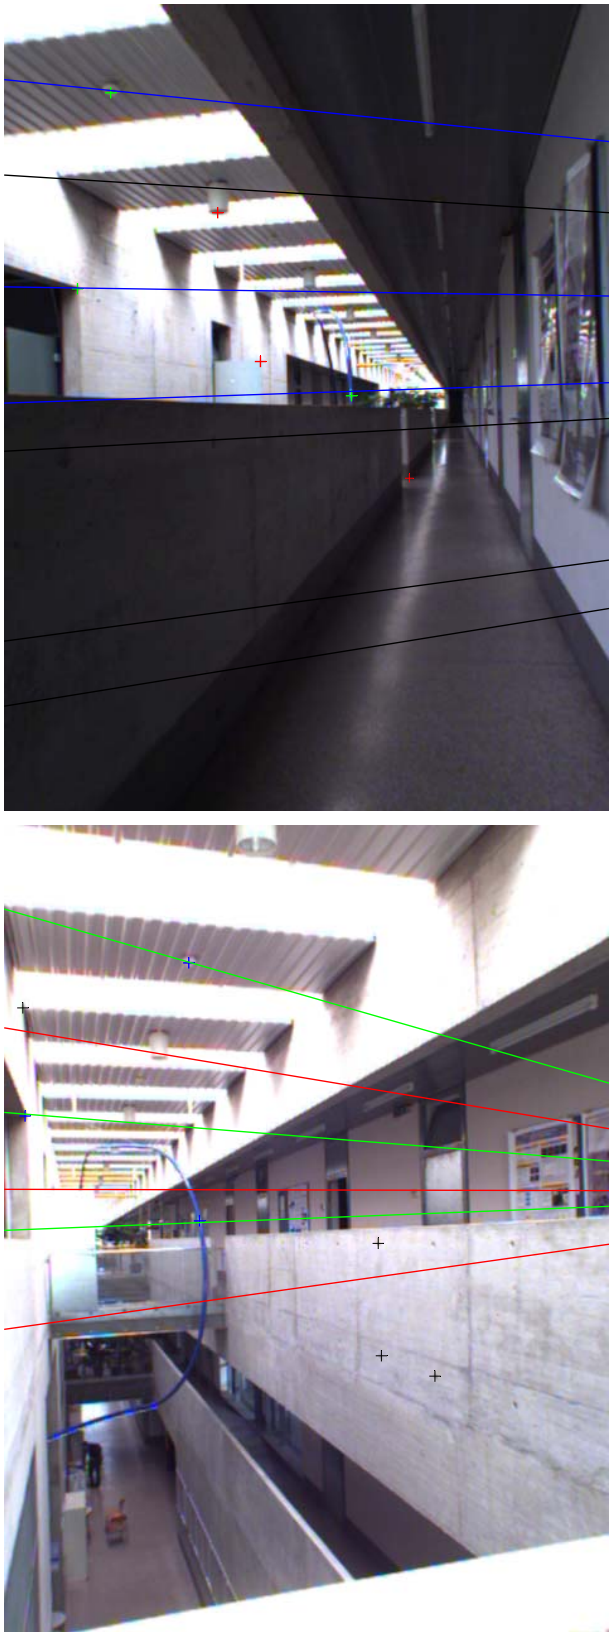

Figure 5: Two views of a corridor scene. The computed epipolar geometry is shown by points and corresponding epipolar lines.