

Efficient Computation of Vanishing Points

Wei Zhang and Jana Kosecka¹

Department of Computer Science

George Mason University, Fairfax, VA 22030, USA

{wzhang2, kosecka}@cs.gmu.edu

Abstract

Man-made environments possess a lot of regularities which simplify otherwise difficult pose estimation and visual reconstruction tasks. The constraints arising from parallel and orthogonal lines and planes can be efficiently exploited at various stages of vision processing pipeline. In this paper we propose an approach for estimation of vanishing points by exploiting the constraints of structured man-made environments, where the majority of lines is aligned with the principal orthogonal directions of the world coordinate frame. We combine efficient image processing techniques used in the line detection and initialization stage with simultaneous grouping and estimation of vanishing directions using expectation maximization (EM) algorithm. Since we assume uncalibrated camera the estimated vanishing points can be used towards partial camera calibration and estimation of the relative orientation of the camera with respect to the scene. The presented approach is computationally efficient and has been verified extensively by experiments.

Key words: Vanishing point estimation, relative orientation, calibration using vanishing points, vision guided mobile and aerial robots.

1 Introduction

The problem of recovering relative orientation of the camera with respect to the scene is of importance in a variety of applications which require some knowledge of the environment's geometry. These range from basic structure and motion or pose recovery problems from single or multiple views, autonomous robotic navigation, manipulation and human computer interaction tasks. Common to these applications is the need for establishment of the relative pose of the camera with respect to the scene reference frame. The desired accuracy varies based on the application. In the case of man-made environments, which possess a lot of regularities, such as presence of sets of parallel and orthogonal lines and planes, the relative orientation can be determined

from vanishing points and lines in the image plane.

This paper describes an approach for the estimation of vanishing points in structured man-made environments. While partial solutions to this problem have been addressed numerous times in the past, the proposed approaches vary in the level of automation of the process, computational complexity, feature detection, assumptions about camera calibration, initialization and grouping stage as well as the choice of line representation and geometric estimation technique. We believe that the robustness and flexibility of the presented approach is superior to the previously suggested techniques and it is also amenable to implementation on robotic platforms. The additional appeal of the technique is in the fact that it does not require calibrated camera. Consequently the estimated vanishing points can be used for partial camera calibration yielding more flexible overall system.

1.1 Paper outline

This paper describes a completely automated process of detecting the vanishing points in the image. We describe the individual steps of the approach as well as the rationale behind the individual choices. The stages include detection of the image line segments, line fitting, initialization, followed by grouping of the lines into dominant vanishing directions and estimation of vanishing points. We assume that the camera is not calibrated and discuss some choices of the objective functions for the vanishing point estimation problem.

2 Related Work

There are numerous articles related to the subproblems and techniques used in our approach. We will review more recent representatives and point out the commonalities and differences between individual stages. The starting point common to all techniques for detection of vanishing points is the line detection and line fitting stage. The traditional textbook approach suggests the edge detection, followed by edge chaining and line fitting. The observation that in man-made environments the majority of line segments is aligned with three principal orthogonal directions can make this stage more efficient. In cases when the camera is cal-

¹This work is supported by NSF grant IIS-0118732 and George Mason University Research Initiative fund.

ibrated, the image line segments are represented as unit vectors on the Gaussian sphere corresponding to the plane normals passing through the center of projection. Several techniques for both grouping and initialization stage on the Gaussian sphere have been proposed [1, 2, 4]. The main advantage of the Gaussian sphere representation is the equal treatment of all possible vanishing directions, including those at infinity, which are represented by vectors parallel to the image plane. The clear appeal of the Gaussian sphere representation is hindered by the fact that it assumes calibrated camera. An alternative image based representations of the line segments have been used in [10, 12, 6].

The techniques proposed for the initialization and grouping stage employ different types of the accumulator space, where the peaks correspond to the dominant clusters of line segments. Both in [1, 4] the Gaussian sphere was suggested as an accumulator space, where the lines correspond to great circles and vanishing directions to the intersections of great circles on the sphere. Hough space has been also used in the accumulation stage [11, 13]. An alternative initialization schemes using the image based representations which considered all pairwise intersections of the detected line segments have been proposed by [12, 2]. While the techniques which consider all pairwise intersections, has been previously shown to lead to more accurate detection they have quadratic running time in the number of line segments. When the number of line segments is large the complexity issues matter, specially in the case when one would like to employ such techniques possibly in an online manner. This point was argued in [1] who suggested linear initialization and grouping technique for the detection of vanishing directions.

At last the previously proposed techniques differ in the type of error measure used to quantify the distance between a line segment and vanishing direction, which determines the objective function to be minimized. For the Gaussian sphere line representations, the objective functions are typically linear and can be minimized in the closed form using weighted linear least squares [4, 1]. In the case of image based representations either the angle between the vanishing line and vanishing direction is minimized or the distance of the line end points from the vanishing direction. In both cases the error measure leads to nonlinear functions of the unknowns [10, 6] and involves iterative optimization techniques.

3 Approach

Our approach is most closely related to the works of [1] and [10]. In the grouping stage we adopt similar strategy to the one suggested by [1] without the need of calibrated cameras and more efficient initialization scheme. We explore different choices of objective functions and demonstrate the fea-

sibility of the Gaussian sphere representation for grouping and estimation in the case of uncalibrated camera. The basic premise of our approach is the observation that the distribution of the gradient orientation reflects the actual orientation of the line segment in the image and that in man-made environments the majority of the line segments is aligned with the three principal orthogonal directions $\mathbf{i}, \mathbf{j}, \mathbf{k}$ associated with the world frame. These observations are used both in the line detection stage and initialization stage and favorably affect the speed of the algorithm.

3.1 Line Detection Stage

The first step of the line detection stage is the computation of image derivatives using, for the efficiency reasons, Sobel edge detector. The line fitting stage follows an approach suggested by [8]. The gradient direction is quantized into a set of k ranges, where all the edge pixels having an orientation within the certain range fall into the corresponding bin and are assigned a particular label. In our case $k = 16$. The edge pixels having the same label are then grouped together using connected components algorithm. For the purpose of calculation of vanishing points we only consider dominant connected components, whose length is beyond certain threshold, which depends on the size of the image (e.g. in our case $t_l = 15$ for 400×300 image size). The line segment candidates are obtained by fitting a line parameterized by an angle θ and distance from the origin ρ :

$$\rho = x \cos(\theta) + y \sin(\theta) \quad (1)$$

Each obtained connected component is a list of edge pixels (x_i, y_i) with the similar gradient orientation, which form the line support regions. The line parameters are directly determined from the eigenvalues λ_1 and λ_2 and eigenvectors \mathbf{v}_1 and \mathbf{v}_2 of the matrix D associated with the line support region:

$$D = \begin{bmatrix} \sum_i \tilde{x}_i^2 & \sum_i \tilde{x}_i \tilde{y}_i \\ \sum_i \tilde{x}_i \tilde{y}_i & \sum_i \tilde{y}_i^2 \end{bmatrix} \quad (2)$$

where $\tilde{x}_i = x_i - \bar{x}$ and $\tilde{y}_i = y_i - \bar{y}$ are the mean corrected pixels coordinates belonging to a particular connected component and $\bar{x} = \frac{1}{n} \sum_i x_i$ and $\bar{y} = \frac{1}{n} \sum_i y_i$. In case of an ideal line one of the eigenvalues should be zero. The quality of the line fit is characterized by the ratio of the two eigenvalues of matrix D , $\nu = \frac{\lambda_1}{\lambda_2}$. The line parameters (ρ, θ) are determined from the eigenvectors $\mathbf{v}_1, \mathbf{v}_2$, where \mathbf{v}_1 is the eigenvector associated with the largest eigenvalue. The line parameters are computed as:

$$\begin{aligned} \theta &= \text{atan2}(\mathbf{v}_1(2), \mathbf{v}_1(1)) \\ \rho &= \bar{x} \cos \theta + \bar{y} \sin \theta \end{aligned} \quad (3)$$

where (\bar{x}, \bar{y}) is the mid-point of the line segment. In practice many detected line segments do not belong to the environment structure and are due to either shadow or shading effects. These spurious line segments are inevitable and depend on the choice of the threshold, which in this stage was

selected globally for the entire image. The threshold choice is later used in computation of the weighting factors for the grouping process.

3.2 Vanishing point estimation

Prior to vanishing point estimation the line segments obtained in the previous stage need to be grouped into the dominant vanishing directions. In general such grouping can be a difficult problem, since any two parallel lines intersect in a vanishing point, possibly yielding a large set of vanishing points. However in the case of man-made environments the dominant vanishing directions are the ones aligned with the principal axes $\mathbf{i}, \mathbf{j}, \mathbf{k}$ of the world coordinate frame and the majority of the parallel lines will be aligned with these directions. We suggest to address the grouping stage and vanishing point estimation stage as a problem of probabilistic inference with an unknown model. In such instances the algorithm of choice is the Expectation Maximization algorithm (EM), which simultaneously estimates the coordinates of vanishing points as well as the probabilities of individual line segments belonging to a particular vanishing directions. The main differences between the previous EM approach [1] and this one, is our assumption of an uncalibrated camera, different likelihood function and more efficient initialization stage.

Our goal is to estimate most likely orthogonal vanishing directions, given a set of line segments, so as to maximize the posterior probability of the vanishing directions. Using Bayes rule, we can express the posterior distribution of the vanishing points given line segments, in terms of the conditional distribution and prior probability of the vanishing points:

$$p(\mathbf{v}_k | \mathbf{l}_i) = \frac{p(\mathbf{l}_i | \mathbf{v}_k)p(\mathbf{v}_k)}{p(\mathbf{l}_i)} \quad (4)$$

Considering the detected set of line segments, majority of them can be classified as arising from one of the principal vanishing directions. Hence for a particular line segment, $p(\mathbf{l}_i)$ can be expressed using the conditional mixture model representation:

$$p(\mathbf{l}_i) = \sum_{k=1}^m p(\mathbf{v}_k)p(\mathbf{l}_i | \mathbf{v}_k) \quad (5)$$

where m is the number of possible vanishing directions. This number will vary depending on the image, but in general we will assume that there are at most four significant models, three corresponding to the dominant vanishing directions and an additional one modeling the outliers process. The lines which do not belong to the vanishing direction aligned with the principal directions are considered to be outliers. The choice of the likelihood term $p(\mathbf{l}_i | \mathbf{v}_k)$ depends on the line representation, as well as the form of the objective function to be minimized.

3.2.1 Line representation: In the perspective camera projection model, the 3D coordinates of points $\mathbf{X} = [X, Y, Z, 1]^T$ are related to their image projections $\mathbf{x} = [x, y, 1]^T$ in a following way:

$$\lambda \mathbf{x} = P g \mathbf{X}$$

where $P = [I_{3 \times 3}, 0] \in \mathbb{R}^{3 \times 4}$ is the projection matrix, $g = (R, T) \in SE(3)$ is a rigid body transformation represented by 4×4 matrix using homogeneous coordinates and λ is the unknown scale corresponding to the depth Z of the point \mathbf{X} . Given two image points \mathbf{x}_1 and \mathbf{x}_2 , the line segment passing through the two endpoints is represented by a plane normal of a plane passing through the center of projection and intersecting the image in a line l , such that $\mathbf{l} = \mathbf{x}_1 \times \mathbf{x}_2 = \widehat{\mathbf{x}}_1 \mathbf{x}_2^1$. The unit vectors corresponding to the plane normals \mathbf{l}_i can be viewed as points on a unit sphere. The vectors \mathbf{l}_i corresponding to the parallel lines in 3D world all lie in the common plane and their vanishing direction corresponds to the normal of that plane. Given two lines the common normal is determined by $\mathbf{v} = \mathbf{l}_1 \times \mathbf{l}_2 = \widehat{\mathbf{l}}_1 \mathbf{l}_2$. Hence given a set of line segments corresponding to parallel lines in 3D, the common vanishing direction \mathbf{v} can be obtained by solving the following linear least squares estimation problem: $\min_{\mathbf{v}} \sum_{i=1}^n (\mathbf{l}_i^T \mathbf{v})^2$. This particular least squares estimation problem has been studied previously in [4] and [7]. Given a set of line segments belonging to the same vanishing direction, the orthogonal least squares solution is applicable regardless of the camera being calibrated. The absence of calibration however affects the grouping stage. In the following section we will demonstrate that the absence of calibration can be remedied by a proper normalization scheme, which does not affect the final estimates.

In the case of uncalibrated camera the image coordinates undergo an additional transformation K which depends on the internal camera parameters:

$$\mathbf{x}' = K \mathbf{x} \quad \text{with} \quad K = \begin{bmatrix} f & \alpha_\theta & o_x \\ 0 & kf & o_y \\ 0 & 0 & 1 \end{bmatrix}$$

where f is the focal length of the camera in pixel units, k is the aspect ratio and $[o_x, o_y, 1]^T$ is the principal point of the camera. Consider now a line in the uncalibrated case specified by the end points \mathbf{x}'_1 and \mathbf{x}'_2 , then:

$$\mathbf{l}'_1 = \widehat{\mathbf{x}}'_1 \mathbf{x}'_2 = \widehat{K \mathbf{x}_1} K \mathbf{x}_2 = K^{-T} (\widehat{\mathbf{x}}_1 \mathbf{x}_2) = K^{-T} \mathbf{l}_1$$

where we have exploited the fact that $A^T \widehat{v} A = \widehat{A^{-1}v}$ for any $v \in \mathbb{R}^3$ and $A \in SL(3)$. Hence in the uncalibrated case the vanishing point \mathbf{v}' computed as a normal to the plane spanned by vectors \mathbf{l}'_1 and \mathbf{l}'_2 ; $\mathbf{v}' = \mathbf{l}'_1 \times \mathbf{l}'_2$ is related to the actual vanishing direction in the calibrated space by the unknown transformation K , namely $\mathbf{v}' = K \mathbf{v}$. In the absence of calibration the distance between two vanishing

¹ $\widehat{\mathbf{x}}$ is a skew symmetric matrix associated with $\mathbf{x} = [x_1, x_2, x_3]^T$.

directions characterized by the inner product between two vectors $\mathbf{v}_1, \mathbf{v}_2$ becomes $\mathbf{v}_1^T \mathbf{v}_2 = \mathbf{v}_1'^T K^{-T} K^{-1} \mathbf{v}_2'$ where $\mathbf{v}_i = K^{-1} \mathbf{v}_i'$. Note that the inner product is now dependent on an unknown matrix $S = K^{-T} K^{-1}$, which causes a *distortion* of the overall space. This observation affects the grouping stage, where we seek so assign the individual line segments to the dominant vanishing directions, which in the calibrated case are well separated on the Gaussian sphere. In order to reduce the effects of the distortion and make the grouping well conditioned we propose to normalize the measurements.

3.2.2 Normalization: Prior starting the estimation process and determining the line segment representation, we first transform all the endpoint measurements by transformation A , in order to make the line segments and vanishing directions well separated on the unit sphere and consequently similar to the calibrated setting:

$$\tilde{\mathbf{x}} = A\mathbf{x}' = \begin{bmatrix} \frac{1}{f^*} & 0 & -\frac{o_x^*}{f^*} \\ 0 & \frac{1}{f^*} & -\frac{o_y^*}{f^*} \\ 0 & 0 & 1 \end{bmatrix} \mathbf{x}'$$

Given an image of size $s = [nrows, ncols]$ the choice of the transformation A is determined by the size of the image and captures the assumption that the optical center is in the center of the image and the aspect ratio $k = 1$ and effectively serves as an approximation of the inverse of the calibration matrix K^{-1} . The focal length in the pixel units is $f^* = nrows$, $o_x^* = \frac{nrows}{2}$ and $o_y^* = \frac{ncols}{2}$. Given the assumptions about optical center and aspect ratio, the chosen focal length f^* is related to the actual focal length by a scale factor. From now on we will drop the $'$ superscript image coordinates and assume all the subsequent computations in this normalized space. Once the final estimates of vanishing points have been computed, the results are transformed back in order to obtain the vanishing point estimates in the actual image coordinates. For each vanishing point $\tilde{\mathbf{v}}_k$ estimated in the normalized space we transform it back in order to obtain the vanishing points in true image coordinates; $\mathbf{v}_k = A^{-1} \tilde{\mathbf{v}}_k$. These can be used further for camera calibration and estimation of relative rotation.

3.3 Expectation Maximization (EM)

Given a set of line segments we would like to find the most likely estimates of vanishing points as well as probabilities of each line belonging to a particular vanishing direction. Our ultimate goal is to estimate the coordinates of all vanishing points so as to maximize the likelihood of the vanishing point estimates given a set of line segments. Given initial estimates of the vanishing points \mathbf{v}_k , the membership probabilities of a line segment \mathbf{l}_i belonging to the k^{th} vanishing direction are computed in the following way:

$$p(\mathbf{v}_k | \mathbf{l}_i) = \frac{p(\mathbf{l}_i | \mathbf{v}_k)p(\mathbf{v}_k)}{\sum_{k=1}^m p(\mathbf{l}_i | \mathbf{v}_k)p(\mathbf{v}_k)} \quad (6)$$

The posterior probability terms $p(\mathbf{v}_k | \mathbf{l}_i)$ represent so called membership probabilities, denoted by w_{ik} and capture the probability of a line segment \mathbf{l}_i belonging to k^{th} vanishing direction. Currently we assume that the prior probabilities of all vanishing directions are equally likely, hence the prior probabilities do not affect the posterior conditional probability. Incorporating the information about prior probabilities of the vanishing directions can provide additional information and make the estimation process better conditioned. The E-step of the EM algorithm amounts to computation of these posterior probabilities, which give us the best guess of the membership probabilities given the currently available vanishing points estimates. The M-step of the algorithm involves maximization of the expected complete log likelihood with respect to the unknown parameters \mathbf{v}_k [9], which yields a maximization of the following objective function:

$$J(\mathbf{v}_k) = \sum_i w_{ik} \log p(\mathbf{l}_i | \mathbf{v}_k) \quad (7)$$

In the noise free case $\mathbf{l}_i^T \mathbf{v}_k = 0$. In the case of noisy estimates we assume that the error represented by $\mathbf{l}_i^T \mathbf{v}$ is a normally distributed random variable with $N(0, \sigma_1^2)$. Then the likelihood term is given as:

$$p(\mathbf{l}_i | \mathbf{v}_k) \propto \exp\left(\frac{-(\mathbf{l}_i^T \mathbf{v}_k)^2}{2\sigma_1^2}\right) \quad (8)$$

The above objective function (equation 7) in case of linear models yields a solution to a weighted least squares problem (one for each model), where each line has an associated weight w_{ik} determined by posterior probabilities computed in the E-step (equation 6). The above objective function in case of linear log-likelihood model yields a solution to a weighted least squares problem; one for each model. In such case the vanishing points are estimated by solving the following linear least-squares problem:

$$J(\mathbf{v}_k) = \min_{\mathbf{v}_k} \sum_i w_{ik} (\mathbf{l}_i^T \mathbf{v}_k)^2 = \min_{\mathbf{v}_k} \|W A \mathbf{v}_k\|^2 \quad (9)$$

Where $W \in \mathbb{R}^{n \times n}$ is a diagonal matrix associated with the weights and rows of $A \in \mathbb{R}^{3 \times n}$ are the detected line segments. Closed form solution corresponds to the eigenvector associated with the smallest eigenvalue of $A^T W^T W A$ and yields the new estimate of \mathbf{v}_k . The EM algorithm is an iterative technique guaranteed to increase the likelihood of the available measurements. We terminate the iteration procedure once the vanishing point estimates reach an equilibrium, *i.e.* $\mathbf{v}^{jT} \mathbf{v}^{(j-1)} < \epsilon_a$. The iterations of the EM algorithm are depicted for different examples in Figure 2 (right). The initially large number of vanishing point estimates is reduced through the merging process to three dominant directions. Once the algorithm converges, we apply the nonlinear refinement step, to achieve a more accurate vanishing point estimate. The nonlinear refinement step computes the Maximum Likelihood Estimate (MLE) of the




Figure 1: The error measure for the nonlinear refinement step is $d^2 + d'^2$.

vanishing point given the line segments, by minimizing the sum of squared distances d, d' of the line end points from the vanishing line (see Figure 1). The nonlinear optimization has been suggested previously by [6] and is especially important for refining the estimates of the vanishing points close to infinity.

In order to account for the line segments which do not belong to any of the vanishing directions we add an additional outlier process. The initial probability of a line segment belonging to the mixture is determined by the highest possible residual angle which a line segment can have with respect to one of the dominant orientations $e_i = \pi/4$. The probability of the outliers is then determined by a Gaussian distribution with a large variance, approximating the uniform distribution.

3.4 EM Initialization

While the EM algorithm is known in each step to increase the likelihood of the measurements given the model, it is often very sensitive to initialization and can converge to a local minimum. Hence it is very important that the algorithm is correctly initialized. The initialization scheme adopted in our approach is based on the observation that the lines that belong to a particular vanishing directions have similar orientation in the image. This is not the case for the line segments whose vanishing point is close to the center of the image plane. In such situation the lines will be initially assigned to a separate groups and merged later in the EM stage. Given a set of detected line segments we form a histogram h_θ of their orientations and search for the peaks in the histogram. The peaks are detected by first computing the curvature $\mathcal{C}(k)$ of the histogram, followed by a search for zero crossings, which separate the dominant peaks. The curvature is computed by subtracting the local histogram mean:

$$\mathcal{C}(k) = h_\theta(k) - \frac{1}{m} \sum_{i=k-\frac{m}{2}}^{k+\frac{m}{2}+1} h_\theta(i) \quad (10)$$

The total number of line orientation bins in our case is $K = 60$ and the size of the integration window $m = 5$. Histogram h_θ is smoothed prior to the curvature computation. The typical number of detected peaks ranges between 2 – 6 and determines the initial number of models in

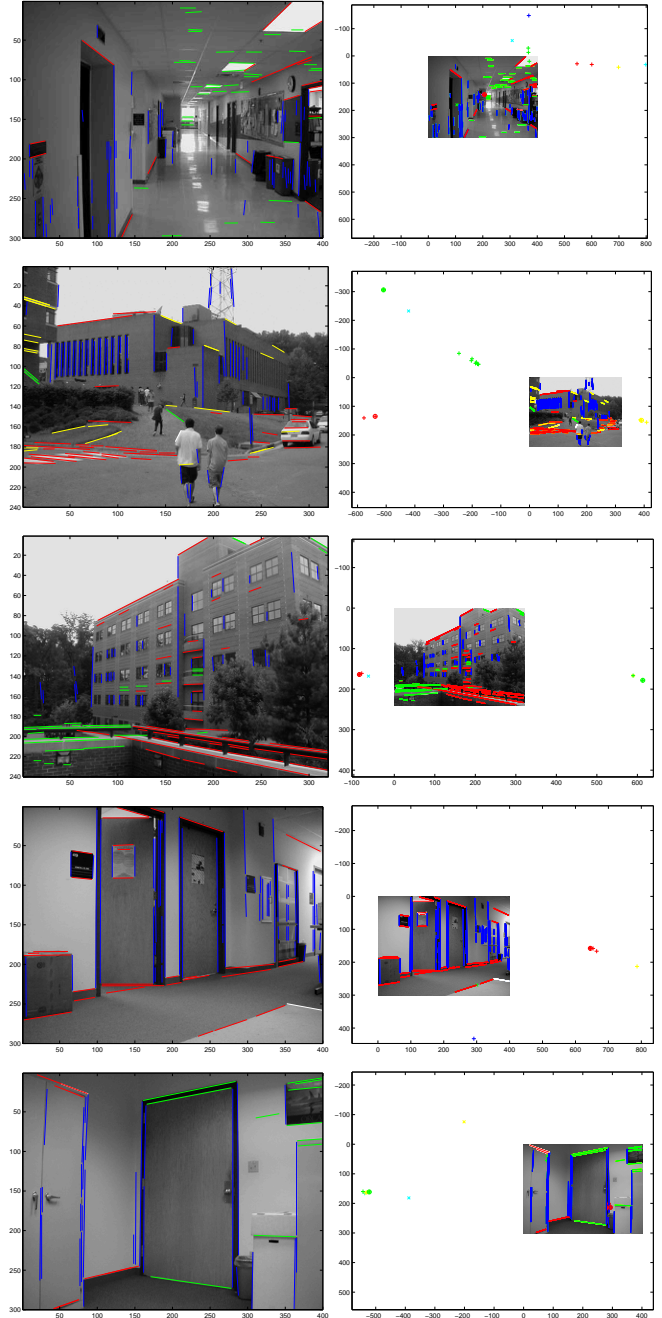


Figure 2: The vertical (horizontal) lines parallel to the image plane yield the vanishing point close to infinity and hence outside of the selected view. The lines are color-coded according to the maximum membership probability and are numbered from top to bottom. The final estimates of vanishing points are reported in Table 1. In case more than three vanishing directions are detected (image 2), the three with the largest number of line segments are retained. In some cases only two vanishing directions are present (image 4).

the mixture formulation. Since the initialization determines more peaks than the number of vanishing directions sought for, this has to be reconciled in the expectation maximization process. The decision to adjust the number of models is made based on the distance between currently available estimates. If the distance between two vanishing points is below some threshold the two mixtures are merged and the number of models is adjusted. During the course of iteration of the EM algorithm, we also adjust the standard deviation of the individual mixtures, to reflect the increasing confidence in the model parameters and memberships of the individual segments.

The performance of the algorithm is reported in Table 1 and demonstrated in Figure 2. More detailed evaluation of the algorithm can be found in [14]. The presented results also depict the iterations of the EM algorithm.

Image No.	Vanishing point estimates		
	\hat{v}_i	\hat{v}_j	\hat{v}_k
1	206.66	-8635.12	188.48
	142.16	491.89	10298.56
2	393.07	-527.08	153.96
	148.94	133.98	-3641.49
3	613.31	-86.13	150.61
	178.13	163.88	-3136.68
4	×	645.19	169.41
	×	157.86	-16339.53
5	293.49	-521.80	362.46
	213.87	161.30	-16575.81

4 Summary and Conclusions

We presented an efficient, completely automated approach for detection of vanishing points from a single view assuming an uncalibrated camera. Along the way the assumptions about the structure of man-made environments, were used towards efficient initialization and grouping stage. In particular it was the fact that the majority of the line segments belongs to one of the dominant vanishing directions aligned with the axes of the world coordinate frame. The estimation and grouping problems were addressed simultaneously using the probabilistic Expectation Maximization algorithm. We are currently refining the algorithm and exploring the use of the prior information for the vanishing point estimation task as well as performance of the technique on a large variety of indoors and outdoors scene with different resolutions. The capability of robustly detecting vanishing points in an automated manner will enable us to employ these types of systems in the context of mobile and aerial robots and facilitate partial calibration of the vision system as well as estimation of relative orientation with respect to the scene. This information has great utility in the context

of basic navigation and exploration tasks in indoor and outdoor environments, where the alternative sensing strategies such as GPS or compass are known to under perform.

References

- [1] M. Antone and S. Teller. Automatic recovery of relative camera rotations for urban scenes. In *IEEE Proceedings of CVPR*, 2000.
- [2] B. Brillaut-O'Mahony. New method for vanishing point detection. *CVGIP: Image Understanding*, 54(2):289–300, September 1991.
- [3] B. Caprile and V. Torre. Using vanishing points for camera calibration. *International Journal of Computer Vision*, 3:127–140, 1990.
- [4] R. Collins. Vanishing point calculation as statistical inference on the unit sphere. In *Proceedings of International Conference on Computer Vision*, pages 400–403, 1990.
- [5] J. Coughlan and A. Yuille. Manhattan world: Compass direction from a single image by bayesian inference. In *IEEE Proceedings International Conference on Computer Vision*, pages 941–7, 1999.
- [6] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [7] K. Kanatani. *Geometric Computation for Machine Vision*. Oxford Science Publications, 1993.
- [8] P. Kahn, L. Kitchen, and E.M. Riseman. A fast line finder for vision-guided robot navigation. *IEEE Transactions on PAMI*, 12(11):1098–1102, 1990.
- [9] G. J. McLachlan and K. E. Basford. *Mixture Models: Inference and Applications*. Marcel Dekker Inc., N.Y., 1988.
- [10] G. McLean and D. Kotturi. Vanishing point detection by line clustering. *IEEE Transactions on PAMI*, 17(11):1090–1095, 1995.
- [11] L. Quan and R. Mohr. Determining perspective structures using hierarchical hough transforms. *P.R. Letters*, 9:279–286, 1989.
- [12] C. Rother. A new approach for vanishing point detection in architectural environments. In *Proceedings of the British Machine Vision Conference*, 2000.
- [13] T. Tuytelaars, M. Proesmans, and L. Van Gool. The cascaded hough transform. In *Proceedings of ICIP*, pages 736–739, 1998.
- [14] W. Zhang and J. Kořecká. Video Compass. George Mason University Technical Report, GMU-TR-2002.