

Illumination Insensitive Robot Self-Localization Using Panoramic Eigenspaces

Gerald Steinbauer¹ and Horst Bischof²

¹ Institute for Software Technology, Graz University of Technology,
Inffeldgasse 16b/II, A-8010 Graz, Austria

² Institute for Computer Graphics and Vision, Graz University of Technology,
Inffeldgasse 16/II, A-8010 Graz, Austria

Abstract. We propose to use a robust method for appearance-based matching that has been shown to be insensitive to illumination and occlusion for robot self-localization. The drawback of this method is that it relies on panoramic images taken in one certain orientation, restricts the heading of the robot throughout navigation or needs additional sensors for orientation, e.g. a compass. To avoid these problems we propose a combination of the appearance-based method with odometry data. We demonstrate the robustness of the proposed self-localization against changes in illumination by experimental results obtained in the RoboCup Middle-Size scenario.

1 Introduction

Mobile robot localization is the problem of determining a robot's pose (i.e. its location and orientation) from sensor data such as odometry, proximity sensors or vision. Self localization is a key problem in autonomous robotics, it has even been referred to as "the most fundamental problem to providing a mobile robot with autonomous capabilities" [1].

Good results have been achieved by using combinations of metric maps of the environment, sensor models, i.e. that model the expected response from the environment, and probabilistic sensor fusion [2, 3]. These approaches show very accurate localization. But generation and storage of the maps and the models are very time and memory consuming. Especially, if the environment gets large.

In contrast, evidence has been provided that topological localization, which recognizes certain spots in the environment, is sufficient to navigate a mobile robot through an environment. This also appears more naturally, if one thinks how humans navigate through a building or a city.

Appearance-based approaches use images to recognize known spots in the environment. This is done by comparing the current image with a set of reference images, previously captured at some reference locations. The approaches differ in the type of the camera used, the representation of the reference images and the calculation of the similarity of different images. Ishiguro et.al. [4] use panoramic cylinder images obtained from an omnidirectional camera and row by

row Fourier transformation as a compact representation. The similarity of images is determined by calculating the sum of the absolute difference of the most significant Fourier coefficients. Menegatti and colleagues [5] extend the method by Monte Carlo Localization to solve the problem of *perceptual aliasing*. That means the response of the current image to several reference images. In [6] topological localization is applied to the RoboCup Middle-Size. A standard camera and an eigenspace based representation are used. The reference images are transformed to eigenimages by Principal Component Analysis and represented by the coefficients in the corresponding eigenspace. The calculation of the similarity of images is done by a k-nearest neighbor algorithm within the eigenspace.

As changes in illumination dramatically effect the appearance of locations in the environment [7], the mentioned methods are sensitive to such changes in illumination. Currently, we have well restricted lighting conditions in the RoboCup Middle-Size, which are a compromise to the vision algorithm applied to robotic soccer today. Due to the decision to introduce a certain amount of natural illumination to the RoboCup Size-Middle and the vision that our robots someday will leave the field and work in more realistic environments, the importance of illumination insensitive algorithms is undoubted.

To cope with such changes in illumination in [8] another representation and calculation of similarity were proposed. An illumination insensitive eigenspace representation and a randomized voting algorithm are used. The illumination insensitive eigenspace approach was originally developed for robust object recognition under varying lighting conditions [9]. The approach exploits the property that the eigenspace representation also holds after linearly filtering the current image and the eigenimages. This filtering is the key to the illumination insensitivity. Additionally, a voting algorithm based on a randomly drawn subset of pixels of the images makes the approach insensitive to highlights, noise and occlusions. The drawback of this method is that it assumes that the reference images and the current image are captured at one certain orientation of the robot. To meet this assumption either the heading of the robot is restricted to that orientation or an additional sensor for orientation, e.g. a compass, is needed. Both are rigorous limitations for the practical use on mobile robots.

In a recent publication [10] a rotation invariant representation of eigenimages was presented. But this representation is not robust, computationally expensive and still sensitive to changes in illumination.

There are also other representations of panoramic images, that are invariant to rotation [11]. But again they are neither robust nor illumination insensitive.

In this work, we propose the combination of the robust illumination insensitive eigenspace approach and sensor fusion with odometry data as a solution to the limitations mentioned above. The basic idea is to use odometry to keep track of the orientation of the robot, to use the predicted orientation to rotate the current image back to the reference orientation and correct the orientation delivered by the odometry by the response of the eigenspace framework.

In the next section we will outline the self-localization based on illumination insensitive eigenspaces. In Section 3 we introduce our extensions and their the-

oretical foundations. In Section 4 we present preliminary results obtained from experiments in the RoboCup Middle-Size scenario. Finally, in Section 5 some conclusions are drawn and future research perspectives explained.

2 Illumination Insensitive Self-Localization Using Eigenspace

In appearance-based localization the robot is provided with a set of views of the environment taken at several locations in the environment. These locations are called reference locations because the robot will refer to them to locate itself in the environment. The corresponding images are called reference images. When the robot moves, it can compare the current view with reference images captured during a training phase. When the robot finds which of the reference image is more similar to the current view it can infer its position in the environment. The problem of finding the position in the environment is reduced to the problem of finding the best match of the current image among the reference images. A higher localization accuracy can be achieved by interpolation between reference points within the eigenspace, while keeping the number of reference images constant. In the remainder of this section we outline the appearance-based approach presented in [8] and its solution for illumination insensitivity.

2.1 Eigenspace Based Recognition

An eigenspace based representation is used for a compact storage of the reference images and a robust calculation of the similarity of images. To build the eigenspace, we first represent the images from the training set as image vectors, from which the mean image is subtracted, $x_i; i = 0 \dots N - 1$, which form an image matrix $X = [x_0 \ x_1 \ \dots \ x_{N-1}]$, $X \in \mathbf{R}^{n \times N}$; where n is the number of pixels in the image and N is the number of images. These training images serve as input for the Principal Components Analysis (PCA) algorithm, which results in a set of p eigenimages e_i , $i = 1, \dots, p$, that span a low-dimensional eigenspace. Eigenimages are selected on the basis of the variance that they represent in the training set. Every original image x_i can be transformed and represented with a set of coefficients $q_{ij} = x_i e_j$, $j = 1, \dots, p$, which represent a point in the eigenspace. That way, every image is approximated as $\tilde{x}_i = \sum_{j=1}^p q_{ij} e_j$. Figure 1 depicts the first four eigenimages for a RoboCup Middle-Size field.

The standard approach to localization is to find the coefficient vector q of the momentary input image y by projecting it onto the eigenspace using the dot product $q_i = \langle y, e_i \rangle$, so that $q = [q_1, \dots, q_p]^T$ is the point in the eigenspace.

If we want the image y to be recognized as its most similar counterpart in the training set (or in a representation constructed by means of interpolation, see [12]), the corresponding coefficients have to lie close together in the eigenspace. However, in the case when the input image is distorted, either due to occlusion, noise or variation in lighting, the coefficient we get by projecting onto the eigenspace can be arbitrarily erroneous [9].

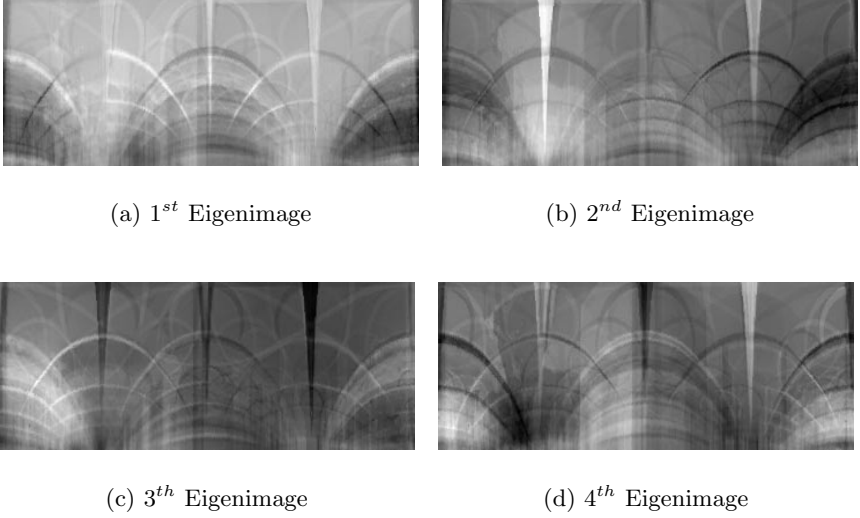


Fig. 1. Eigenspace representation of reference images (Experiments on a RoboCup Middle-Size field)

However, it has been shown, that one can also calculate the coefficient vector q by solving a system of k linear equations on $k \geq p$ points $r = (r_1, \dots, r_k)$

$$y_{r_i} = \sum_{j=1}^p q_j e_{j r_i} \quad 1 \leq i \leq k \quad (1)$$

using a robust equation solver and multiple hypotheses [13].

In [14] it has been shown how this method can be used to allow robust localization in presence of occlusions. However it does not solve the problem of illumination.

2.2 Illumination Insensitivity

The method presented in [15] takes the computations of parameters one step further. Since Eq. (1) is linear, it also holds that $(f * x)(r) = \sum_{i=1}^p q_i (f * e_i)(r)$, where f denotes a filter kernel. This means that if we convolve both sides of the equation with a filter kernel, the equality still holds. Therefore, we can calculate the coefficients q_i also from the filtered eigenimages if we filter the input image.

By using a set of t linear filters \mathcal{F} we can construct a system of equations

$$(f_s * x)(r) = \sum_{i=1}^p q_i (f_s * e_i)(r) \quad s = 1, \dots, t. \quad (2)$$

It is now possible to calculate the coefficients q either by using k points, or using t filter responses at that single point, or a combination of these two.

It is well known from the literature that gradient-based filters are insensitive to illumination variations. By taking a filter bank of gradient filters in several orientations, we can therefore augment the descriptive power of the representation and achieve illumination invariance in the recognition phase.

Illumination invariant localization of a mobile robot can therefore be performed as follows: once the eigenspace is built, we filter the eigenimages by a bank of filters. Then, for localization, the momentary input image y from the panoramic camera has to be filtered with the same filters; only after that we retrieve its coefficient vector q using the robust equation solver. The calculated coefficients are used to infer the momentary location of the robot.

2.3 Robust Voting

In order to robustly recover the coefficients in the presence of noise and occlusions a robust voting algorithm is used. The voting consists of the following steps: Multiple hypotheses are generated by robustly solving the sets of equations obtained from a random selection of pixel subsets. Then the nearest neighbor of each hypothesis in the eigenspace (coefficient vector of the closest point on the parametric manifold) is determined. This selected coefficient vector gets a vote. For voting we use a voting function $v(d) \rightarrow [0, 1]$ which gives votes that are inversely proportional to the distance d of the hypothesis from its nearest neighbor. Based on the distribution of the votes we decide whether the coefficients should be accepted (e.g. coefficients with accumulated votes above a given threshold). We use the following voting function:

$$v(d_{ij}, \sigma) = e^{-\frac{d_{ij}^2}{2\sigma^2}}, \quad (3)$$

where

$$d_{ij} = \arccos \frac{a_i^T a_j^{(t)}}{\|a_i\| \|a_j^{(t)}\|} \quad (4)$$

is the angle between the estimated coefficient vector a_i and the nearest coefficient vector in the eigenspace $a_j^{(t)}$. We use the angle instead the Euclidean distance between the coefficients as criteria, because coefficient vectors with the same direction but different lengths represent the same image but in different brightness. The parameter σ determines the width of the voting function.

3 Keeping Track of the Orientation

The approach presented in Section 2 has been shown to be robust against changes in illumination, noise and occlusion. The drawback of that approach is that it relies on reference and current images captured in one certain orientation or that the orientation is known through an additional sensor, e.g. a compass. The reason for that limitation is that the eigenspace-based representation is not invariant to rotation. Therefore, only images taken in the same orientation as the reference

images can be recognized in a robust way. Instead of capturing images at one orientation one could take images at several orientations (e.g. all 10°), but this approach has several drawbacks: First, the number of images and therefore the training time is significantly increased. Second, and more important due to the larger set of images which need to be represented, the number of eigenimages needs to be increased, which on the one hand increases the running time and more importantly as experiments have demonstrated decreases the robustness of the whole approach.

In order to overcome this limitation we use a combination of the robust eigenspace-based approach and sensor fusion with odometry data. The basic idea is to use a Kalman filter to keep track of the orientation of the robot. The odometry data are fused into the filter and provide a prediction of the orientation Θ of the robot. This prediction is used to rotate the current image back into the orientation the reference images were captured. This rotated image is used for the localization step. But, instead of using only one rotated image we repeat the localization step with a set of images rotated by angles drawn from the predictive distribution. That image which gets the highest response from the recognition process determines the new position of the robot. Its corresponding rotation is used to correct the Kalman filter.

3.1 Sensor Fusion

Sensor fusion of odometry data and data from other sensors (e.g., vision, proximity sensors) with Kalman filters is a commonly applied method. So we skip a deeper discussion of this topic. An overview on Kalman filters could be found in [16]. But it should be mentioned that we use the fusion method presented in [17]. The method is an extension of the standard Kalman filter. It uses a bank of Kalman filters working on the same state vector to deal with asynchronous measurements. This is a common problem in real robot systems, as also in our system.

3.2 The Extended Localization

In [18] and [19] it has been shown that the distance between the coefficient of a rotated image and the coefficient of its reference image smoothly increase with the absolute rotation angle, while still keeping the shortest distance to its reference coefficient in the eigenspace, assuming moderate rotation. This property is used in the development of an extension to the approach presented in Section 2 which preserves the illumination insensitivity and overcomes the limitations caused by rotated images.

The extended localization algorithm can be outlined as follows: In a first step N omnidirectional reference images are captured around the environment in one certain reference orientation, i.e. $\Theta_R = 0$. The reference images are unwrapped creating panoramic cylinders (see Figure 2). Panoramic cylinders have the advantage that rotating the original omnidirectional image is only a row wise shift on the panoramic cylinder. Using the algorithm of section 2 the set of panoramic cylinders is transferred into the filtered eigenspace representation.




Fig. 2. An omnidirectional image (left) and its panoramic cylinder (right)

The localization cycle consists of the following steps: We assume that the initial orientation is known. Odometry data are fused into a Kalman filter at the time they are available. The Kalman filter provides a prediction for the true orientation of the robot and its uncertainty about this orientation Θ_t at time t , $KF(t) \rightarrow \langle \hat{\Theta}_t, \hat{\sigma}_t \rangle$, where $\langle \hat{\Theta}_t, \hat{\sigma}_t \rangle$ determines a normal distribution for the true orientation Θ_t of the robot. When a new omnidirectional image is captured at time t the image is transferred into its panoramic cylinder and a set of orientation hypotheses $\Sigma_t = [\Theta_{t,1}, \dots, \Theta_{t,M}]$ are randomly drawn from the distribution $\langle \hat{\Theta}_t, \hat{\sigma}_t \rangle$. With high probability Σ_t contains the true orientation Θ_t of the robot.

Instead of performing the recognition step only once for the captured image, the recognition step is repeated M times on the captured image shifted by $-\Theta_{t,i}$, $i = 1, \dots, M$. This generates M votes for all N reference images: v_{ij} ; $i = 1, \dots, M$; $j = 1, \dots, N$. The shifted images are equivalent to images captured around the reference orientation Θ_R . The property mentioned at the beginning of this section guarantees that the shifted image i with the rotation closest to the true orientation of the robot Θ_t , $i = \operatorname{argmin}_j (|\Theta_{t,j} - \Theta_t|)$, will collect the highest votes v_{ij} for its corresponding reference image j . The corresponding reference position of the reference image j is reported as the new determined location of the robot. The corresponding orientation Θ_i of the captured image is used as a measurement update for the Kalman filter.

4 Experimental Results

We evaluated the extended Localization by carrying out several real and simulated experiments in the RoboCup Middle-Size scenario. The experiments were conducted using robots of our RoboCup Middle-Size team [20]. Our Middle-Size test-field served as the test environment. The test-field is a rule-compliant field except of its size (5m x 6m). It is situated under a glass roof in an open hall. Therefore, the field is directly effected by changes of illumination during the day. The change of the appearance of a location on the field is illustrated in Figure 3.




Fig. 3. The appearance of the same location at noon (left), in the afternoon (center) and in the evening (right)

4.1 Real Experiments

The real experiments were carried out in three runs. In a first step the field was divided into a grid with 1 m resolution. The points on the grid served as reference locations on the field. This led to a total number of 42 locations. At noon for all reference locations a reference image in the reference orientation $\Theta_R = 0$ was captured. All reference images were transformed into a panoramic cylinder with a resolution of 360 x 145 pixels. Using the method described in Section 2 the reference images were transformed into the filtered eigenspace representation. For representation only 15 of the 42 available eigenimages are enough. Due to performance reasons we used a bank of 3 gradient filters with a 3x3 filter kernel instead of the recommended steerable filters.

Within the experiments we used $M = 5$ image orientations randomly drawn from the prediction of the Kalman filter $\langle \hat{\Theta}_t, \hat{\sigma}_t \rangle$ and tested 50 hypotheses with 1000 randomly selected pixel for each rotated image. The σ for the voting algorithm was set to 0.3. The measurement noise of the odometry data was determined during ground truth measurements [21]. The level of measurement noise for the angle measurement was optimized by hand and was kept constant during the experiments.

Three test runs were conducted at noon (immediately after the reference images were captured), in the afternoon and in the evening. Prior to each run the illumination on the field was measured to document the changes in illumination during the runs and its unsteadiness across the field. The illumination was measured on the 42 reference locations and the minimum, the maximum, the mean and the standard deviation was calculated. The results are shown in Table 1.

Table 1. Change of illumination between the test runs

Run	Min/Lux	Max/Lux	Mean/Lux	Std.Dev/Lux
Reference/Noon	335	1288	690.9	248.8
Afternoon	219	857	498.6	157.9
Evening	111	320	203.7	63.1

Table 2. Results of the Localization under different illumination

Test Run	Trials	Reference Positions passed	Reference Positions recognized	Correct Recognitions/%	Tracks Lost
Noon	10	225	225	100	0
Afternoon	10	231	220	95	0
Evening	10	193	161	83	2

Note that the parameters of the camera remained unchanged during capturing the reference images and all runs and no artificial illumination was used.

The test runs were conducted as follows: For each run (noon, afternoon and evening) the robot was randomly placed ten times on a reference location with a random but known orientation. The robot then randomly moved around for a while. Each time the robot passed a reference location (robot entered a cycle of 30 cm diameter around the reference location) it was recorded if the reference location was correctly recognized. The results shown in Table 2 document clearly the performance of the extended localization under different illumination.

In the noon and afternoon runs the robot correctly recognized nearly all reference locations and never lost track of the orientation. The latter is crucial for continuing correct recognition of reference locations. In the evening run the recognition ratio decreased and the robot twice lost track of the orientation due to the very bad illumination (see Figure 3). The lost track was caused by wrong votings on the shifted images and their corresponding rotation, which prevented an adequate correction of the odometry data. Note that we used the Matlab to C++ compiler to convert the Matlab-prototype into a C++ module, executable on the robot. That lead to a frame-rate of only 0.8 Hz. We suppose that a speedup of the implementation will decrease the number of lost tracks, due to the faster correction of erroneous odometry data.

4.2 Offline Experiments

To verify the smooth decreasing of votes for rotated images and the robustness of the method against noise and occlusion we conducted two offline experiments. First we offline calculated the votes for the rotated version of an image captured near a reference location. Figure 4 shows the votes for an image captured in the reference orientation Θ_R near reference location 20. It clearly shows the smooth decreasing of the votes for reference location 20 with increasing absolute rotation angle. For an absolute angle approximately below 10° a clear voting for the reference image 20 is provided. This is an encouraging result, as predictions of the orientation with a higher error do not occur very frequently.

In the second offline experiment we evaluated the robustness of the method against noise. Therefore, we used the images and orientations recorded in the real run at noon. We did the same experiment as in the real run except that we introduced a certain amount of noise into the recorded images. This is done by randomly replacing a certain percentage of the pixels in the image by random pixels. Figure 5 depicts an image before and after introducing 50% random pixels.




Fig. 4. Collected votes for an image take near reference position 20 in respect to shifts of the image




Fig. 5. Original image (a) and the same image after introducing 50% random pixels (b)




Fig. 6. Recognition rate of reference locations in respect to the amount of introduced random pixels




Fig. 7. Original image (a) and the same image after introducing 70% occlusion (b)




Fig. 8. Recognition rate of reference locations in respect to the amount of introduced occlusion (1000 sampled pixel and 50 hypotheses)

The results are shown in Figure 6. It shows clearly the robustness of the method against noise. Note, due to performance reasons we used a very small number of hypotheses (50) and selected pixels (1000) during the previous experiments. Anyway, the recognition rate remained satisfyingly high until an amount of approximately 25% of random pixels. When we increase the number of hypotheses or the number of selected pixels the robustness against noise increases significantly. But of course also the computational costs raises.

In the third offline experiment we evaluated the robustness of the method against occlusion. The third experiment was the same as the prior experiment except that we introduced a certain amount of occlusion instead of noise. This was done by inserting black bars with different width into the images. We used 1000 sampled pixels and tested 50 hypotheses. Figure 7 depicts an image before and after introducing 70% occlusion.

The results are shown in Figure 8. It shows clearly the robustness of the method against occlusion. Up to 70 % of occlusion the method is robust against occlusion and all positions were recognized correctly.

5 Conclusion and Future Work

In this work we propose an extension of the appearance-based self-localization presented in [8] which overcomes the limitations of that method. The limitations are restrictions of the heading of the robot to a reference orientation or the need of an additional sensor for orientation. The extension is based on a combination of the illumination insensitive localization and sensor fusion with odometry data. A prediction for orientation is used for a robust recognition process and its rotation sensitive feedback is used to correct the odometry data. Preliminary experiments within the RoboCup Middle-Size scenario show that the proposed extended localization is robust against changes of illumination, noise and occlusions while it overcomes the limitations mentioned above. However, a speedup of the implementation is necessary to reliably keep track of the orientation under very bad illumination.

Furthermore, we believe that a combination of the illumination insensitive localization with particle filter methods will further improve the method. Especially, because the robot is unable to recover from dramatic errors in the orientation and it has to know its initial orientation.

Currently, we are working on the application of our extended localization in larger environments, e.g. office buildings. As the recording of reference images and its location by hand is time consuming in large environments, we perform research on map-building and SLAM in combination with the illumination insensitive localization.

Acknowledgements

The research presented in this paper has been supported by Land Steiermark under grand 40Ro03-PE "RoboCup 2004". The authors gratefully acknowledge the support of Knapp Automation & Logistik, Sick AG, Kapsch, Saft, Siemens, jumpTec, Farnell In One and BTI for setting up the Mostly Harmless team.

References

1. I.J. Cox. Blanche - An Experiment in Guidance and Navigation of an Autonomous Robot Vehicle. *IEEE Transactions on Robotics and Automation*, 7(2):193–204, 1991.
2. Dieter Fox, Wolfram Burgard, Frank Dellaert, and Sebastian Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *AAAI/IAAI*, pages 343–349, 1999.
3. F. Dellaert, W. Burgard, D. Fox, and S. Thrun. Using the condensation algorithm for robust, vision-based mobile robot localization. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, Fort Collins, CO*. IEEE, 1999.
4. H. Ishiguro and S. Tsuji. Image-based memory of environment. In *Proc. of Int. Conf. on Intelligent Robots and Systems (IROS'96)*, pages 634–639, 1996.

5. E. Menegatti, M. Zoccarato, E. Pagello, and H. Ishiguro. Image-based Monte-Carlo Localisation without a Map. In *Proc. of the 8th Conference of the Italian Association for Artificial Intelligence (AI*IA) Pisa - Italy*, 2003.
6. Goncalo Neto, Hugo Costelha, and Pedro Lima. Topological navigation in configuration space applied to soccer robots. In *RoboCup 2003 International Symposium, Padova, Italy*, 2003.
7. Gerd Mayer, Hans Utz, and Gerhard Kraetzschmar. Playing Robot Soccer under Natural Light: A Case Study. In *RoboCup 2003 International Symposium Padua, Italy*, 2003.
8. M. Jogan, A. Leonardis, H. Wildenauer, and H. Bischof. Mobile robot localization under varying illumination. In *Proc. of 16th International Conference on Pattern Recognition*, 2002.
9. Horst Bischof, Horst Wildenauer, and Ales Leonardis. Illumination Insensitive Recognition using Eigenspaces. *Computer Vision and Image Understanding (to appear)*, 2004.
10. Matjaz Jogan and Ales Leonardis. Robust localization using an omnidirectional appearance-based subspace model of environment. *Robotics and Autonomous Systems*, 45(1):51–72, 2003.
11. Tomas Pajdla and Vaclav Hlavac. Zero phase representation of panoramic images for image based localization. In *8th International Conference on Computer Analysis of Images and Patterns*, pages 550–557, 1999.
12. S. K. Nayar, S. A. Nene, and H. Murase. Subspace methods for robot vision. *IEEE Transaction on Robotics and Automation*, 12(5):750 – 758, 1996.
13. Ale Leonardis and Horst Bischof. Robust recognition using eigenimages. *Computer Vision and Image Understanding: CVIU*, 78(1):99–118, 2000.
14. M. Jogan and A. Leonardis. Robust localization using panoramic view-based recognition. In *15th ICPR*, volume 4, pages 136–139. IEEE Computer Society, 2000.
15. H. Bischof, H. Wildenauer, and A. Leonardis. Illumination insensitive eigenspaces. In *Proc. Int. Conf. on Computer Vision*, volume 1, pages 233–238, 2001.
16. P. S. Maybeck. The Kalman filter, An introduction to concepts. In *Autonomous Robot Vehicles*, pages 194–204, 1990.
17. L. Drolet, F. Michaud, and J. Cote. Adaptable sensor fusion using multiple kalman filters. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2000.
18. H. Murase and S. K. Nayar. Illumination planning for object recognition using parametric eigenspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(12):1219–1227, 1994.
19. Ales Leonardis Matjaz Jogan, Emil Zagar. Karhunen-loeve transform of a set of rotated templates. *IEEE Trans. on Image Processing*, 12(7):817–825, 2003.
20. Gerald Steinbauer, Michael Faschinger, Gordon Fraser, Arndt Mühlenfeld, Stefan Richter, Gernot Wöber, and Jürgen Wolf. Mostly Harmless Team Description. In *Proceedings of the International RoboCup Symposium*, 2003.
21. Lorenz Mayerhofer. Odometry correction for the keksi omnidrive using an echo state network. Technical report, Institute for Theoretical Computer Science University of Technology Graz, 2003.