

Robust localization using an omnidirectional appearance-based subspace model of environment

Matjaž Jogan*, Aleš Leonardis

Computer Vision Laboratory, Faculty of Computer and Information Science, University of Ljubljana, Tržaška 25, SI-1001 Ljubljana, Slovenia

Received 25 February 2002; received in revised form 10 March 2003

Communicated by F.C.A. Groen

Abstract

Appearance-based visual learning and recognition techniques that are based on models derived from a training set of 2D images are being widely used in computer vision applications. In robotics, they have received most attention in visual servoing and navigation. In this paper we discuss a framework for visual self-localization of mobile robots using a parametric model built from panoramic snapshots of the environment. In particular, we propose solutions to the problems related to robustness against occlusions and invariance to the rotation of the sensor. Our principal contribution is an “*eigenspace of spinning-images*”, i.e., a model of the environment which successfully exploits some of the specific properties of panoramic images in order to efficiently calculate the optimal subspace in terms of principal components analysis (PCA) of a set of training snapshots without actually decomposing the covariance matrix. By integrating a robust recover-and-select algorithm for the computation of image parameters we achieve reliable localization even in the case when the input images are partly occluded or noisy. In this way, the robot is capable of localizing itself in realistic environments.

© 2003 Elsevier Science B.V. All rights reserved.

Keywords: Self-localization; Panoramic eigenspace; Robustness; Appearance-based model

1. Introduction

With the rapid developments in the field of computer vision, more and more robot systems use cameras that provide them with pictorial information about their surroundings. Compared to other sensors like sonars, laser beams or infrared cells, cameras are conveying a much richer visual information on the world around us, providing us with cues about texture, shape, and illumination. However, as a consequence, the task of interpreting this information in a useful way is very complex.

Self-localization is one of the crucial capabilities of a mobile robot that enables autonomous navigation in the environment. While estimating the distance to some object can efficiently be performed by sonars or laser beams, visual systems were designed that can construct a 3D model of the surroundings by using only the visual input, and then exploit that model for mobile robot self-localization. Recent examples of such systems are given by Se et al. [27] who use scale-invariant image features as landmarks to build a 3D map and estimate robot ego-motion by matching them, or by Murray and Jennings [21], who use a stereo camera rig to produce a 2D occupancy map.

Over the last decade, several researchers have questioned, whether models obtained by a 3D reconstruction of the environment are really indispensable

* Corresponding author.

E-mail addresses: matjaz.jogan@fri.uni-lj.si (M. Jogan), alesl@fri.uni-lj.si (A. Leonardis).

for a successful artificial vision system. By following the findings of psychologists, more and more theories awoke that describe the visual process as a task of recognizing and associating 2D images. In computer vision, this resulted in a series of successful algorithms that use 2D images of objects to build a representation based solely on the appearance itself [16,20,22,23]. In practice, this means that the representation does not contain explicit information on 3D metrics, but is rather constructed in a way that allows for view recognition using algorithms that have roots in statistical analysis of patterns. Such appearance-based models are expected to have a plethora of advantages over the 3D models. Instructive collections of work that stress the relation between the 3D reconstruction approach and the appearance-based visual recognition can be found in [1,26,29].

In robotics, appearance-based learning and recognition techniques have received most attention in visual servoing and self-localization, which can both be used in autonomous navigation. By defining the task of self-localization as the problem of recognizing the momentary view of the environment and then associating it with the position of the robot in the physical space, it is evident that appearance-based models provide the appropriate framework. This brings several advantages—firstly, there is no need for intensive preprocessing of the visual input, as there is no need to search for prominent features and their correspondences.

On the contrary, when using an appearance-based model, the image can be used without extracting local features, as visual information gets interpreted in its basic form. Furthermore, 3D models are hard to build for complex environments.

Several research groups have presented successful methods for appearance-based localization that differ mostly in the type of optical sensors and in the way the models are built [2,7,9,11,13,15,24]. We will give a brief overview of some of these approaches in Section 2.

Additional motivation for this research comes from the discoveries on navigation strategies of insects, that are, although limited in the brain size, capable of amazingly confident self-localization. In fact, some studies (see [14] and references therein) imply that wood ants may be using a representation of the environment that is built from wide-angle snapshots of the

scene. Localization is then performed by comparing the instantaneous view with the stored snapshots. According to this and some other studies, e.g. by Dill et al. [6], the patterns are processed retinotopically, i.e., the snapshot is not segmented, but interpreted as a whole.

The notion that ants can localize themselves using methods that resemble appearance matching although their processing resources are rather limited is a sufficient reason for looking further in solutions that nature itself is serving us. The next useful hint is that having a wide field-of-view can be helpful for localization. By using a catadioptric camera [8] which acquires panoramic images, we can extend the robot's field-of-view far beyond that provided by classical cameras. Furthermore, we can expect that such views taken from nearby positions and oriented in the same direction will be strongly correlated, as the visual information contained in the panoramic images, does not change so abruptly as it does on the image-plane of standard cameras with a limited field-of-view. This allows us not only to design an efficient strategy based on the distance in the image space, but also to build a compact representation that eliminates redundancy.

In our work we describe a method for localization of a mobile system (Fig. 1) which is based on recognizing a panoramic view using a model built from 2D panoramic images acquired in the learning phase. The model is constructed by using an approximation of the image set, which gets represented by a low-dimensional set of principal components or *eigenvectors* that span a low-dimensional *eigenspace*.




Fig. 1. Mobile robot with a mounted catadioptric panoramic sensor.

Projections of images on this subspace are the basis for calculating a denser interpolated hypersurface. Localization is subsequently performed by estimating the projection of the momentary panoramic robot's view in the eigenspace and by searching for the best match on the hypersurface.

A severe limitation of the standard projection method is the sensitivity to noise and occlusion in the process of recognition. This is because the momentary input image is usually used as a complete pattern when calculating the projections. If the image is occluded or noisy, the projections calculated will be far from the correct result. To overcome this, a robust algorithm for the retrieval of parameters will be designed.

Another important question when using panoramic images is how to distinguish between images that are taken under different orientations of the robot. Such images have the same pictorial content, yet they are rotated for an angle (phase) of ϕ . We discuss that the approaches used until now fail to simultaneously achieve robustness and insensitivity to in-plane rotation.

As the principal contribution, we propose a specific eigenspace representation, called “*eigenspace of spinning-images*”, which makes localization simultaneously robust and insensitive to in-plane orientation of the robot. The representation exploits the fact that a set of rotated templates carrying identical pictorial information can be compressed in the eigenspace in an efficient way. Further, we show that a significant compression ratio can be achieved even when integrating a number of rotated images that all represent a single location. By doing so, we also have the advantage of being able to match the incoming image directly to the whole set of rotated images in the recognition phase (Fig. 2).

The paper is organized as follows. We start with a review of the related work and, in Section 3, continue with the principles of localization using a panoramic eigenspace approach stressing its advantages and drawbacks. In Section 4 we describe how to efficiently perform the learning phase, where we construct a set of basis images which span an eigenspace that describes all the possible orientations of images in a compact and fast-to-compute fashion. In this way, we obtain the “*eigenspace of spinning-images*”, which we use as the base model for the estimation of the momentary position of the robot in the localization phase.




Fig. 2. The eigenspace of spinning-images allows to match the novel input image to a collection of images in multiple orientations.

In Section 5 we show how it is possible to achieve good localization even in the presence of noise and occlusion by introducing a robust method for projecting the images on the eigenspace. Finally, in Section 6 we give the results of the experiments of localization in several indoor environments. We conclude with a summary and an outline of our ongoing work.

2. Related work

Back in 1996, Ishiguro and Tsuji [9] used 2D images for localization of a mobile visual sensor. In order to compress the images, they use the subspace of Fourier harmonics, i.e., they calculate the Fourier coefficients to represent images in a lower-dimensional subspace. They refer to the appearance-based model as to “the iconic memory of scenes”. They already employ a panoramic sensor and warp the images to a cylindrical representation. The coefficients indicate the features of global views at the reference points and the phase components indicate the orientation of the views. The experiments were made in an office environment, which was represented by a map based on an energy spring model. However, they comprise only place recognition by indexing and do not try to generate an interpolated model. The representation used is not robust, since the Fourier transform is inherently a non-robust transformation. An occlusion in the image influences the frequency spectra in a non-predictable way and therefore arbitrarily changes the coefficients arbitrary.

An interesting algorithm which is based on the encoding of non-panoramic images in the parametric eigenspace was proposed by Maeda et al. [18]. Recognition is based on an active vision strategy which helps to resolve the ambiguities of eigenspace analysis by deducing the position from several neighboring views.

Aihara et al. [2] were probably the first to use the panoramic eigenspace approach. In order to avoid the problem of rotation of the sensor around the optical axis, they use row-autocorrelated transforms of cylindrical panoramic images. They tested the system first in an indoor corridor and then by mounting the camera on a rooftop of a van and making a tour of their campus. The approach suffers from less accurate results for images acquired on novel positions, since by autocorrelating the images some of the information is lost. Moreover, the process of autocorrelating the image is non-robust, meaning that any occlusion in the image may result in an erroneous localization.

In our first report on experiments with panoramic eigenspaces of unprocessed images [11] we suggested that alternative representations have to be found in order to handle the problem of in-plane orientation of the sensor. Independently, Pajdla and Hlaváč [24] proposed to estimate a reference orientation from images alone with the zero phase representation (ZPR). ZPR, in contrast to autocorrelation, tends to preserve the original image content while at the same time achieving rotational independence, as it orients images by zeroing the phase of the first harmonic of the Fourier transform of the image. The experiments indicate that images taken at nearby positions tend to have the same reference orientation. The method is, however, sensitive to variations in the scene, since it operates only with a single frequency using a global transform.

Both Aihara and Pajdla tried to overcome the problem of multiple orientations by applying a global transform to the image. It is, however, well known that the eigenspace methods are inherently not robust when global transformations of images are used [16]. Imagine that we obtain an image of the environment in the absence of people and one that is being partially occluded by persons present in the room. Either autocorrelation or Fourier transform of both images would take quite different values.

Recently, Paletta et al. [25] robustified the panoramic eigenspace approach by applying a Bayesian reasoning over local image appearances,

which helps in rejecting false hypotheses that do not fit the structural constraints of the corresponding spline. They achieved robustness and handled rotation by dividing the panoramic image into overlapping vertical strips representing unidirectional camera views. Distributions of sector images in eigenspace are represented by mixture of Gaussians to provide a posterior distribution over potential locations. However, the local windowing introduces ambiguities that have later to be resolved through Bayesian probabilistic framework.

Gaspar et al. [7] used panoramic images unwrapped to a birds eye view. This is useful for orienting the images in an indoor environment such as a corridor, since it can be aligned using straight lines. They claim that a wide field-of-view provides enough stability for the basic principal components analysis (PCA) algorithm.

Vlassis et al. [31] augmented the learning phase by a supervised technique which resulted in an optimized feature space. In [15] they use a kernel smoothing model for the interpolation and study the effect of the number of features on localization. However, they do not address the problems of occlusion and rotation that are of our primary interest.

From this review of the related work we can conclude that the problems of rotation and robustness were thoroughly discussed only in the work of Paletta et al. [25]. While this method proved to be successful, the method that we are going to introduce solves the problem on the level of the representation, as it provides a generalized representation of rotation and applies the basic parametric eigenspace recognition schema. The main advantage of our “*eigenspace of spinning-images*” representation is that it preserves the appearance while solving the problem of rotation. The issue of robustness is addressed by a modified algorithm for the calculation of the parameters of the novel input image which is based on the selection of subset of points that are not part of occluded regions.

3. The principles of localization using a panoramic eigenspace

For building an appearance-based model from a set of images, the eigenspace approach proved itself as a viable one. This approach is a twofold procedure: in the *learning stage*, several panoramic snapshots are

acquired which together form a good depiction of the environment. These images are taken when the robot is exploring the environment either by following some training strategy or under the supervision of a human operator. Then, an orthonormal basis is computed, which allows us to represent every image of the training set with only a few parameters. PCA offers means for calculating an optimal basis using singular value decomposition (SVD) of the covariance matrix, representing the images in a low-dimensional subspace that is an optimal linear approximation of the original set in the least squares sense. This allows for efficient memory consumption, efficient matching, and localization. The training images have to be collected first in order to build the eigenspace. Alternatively, one can employ an incremental algorithm, where images are being sequentially integrated in the model [3].

Let us assume we take images on a regular grid with an appropriate density. The panoramic images are usually taken with a catadioptric camera which produces round images like the one depicted in Fig. 3(a). In practice, we warp these images to a cylindrical representation like the one depicted in Fig. 3(b). The reason to do so is that the rotation of a panoramic image

translates to a row wise shift in a cylindrical image. It is therefore possible to virtually simulate the in-plane rotation of the camera by just shifting the rows.

When all of the images are collected by the robot, we calculate the eigenspace which efficiently describes the acquired data. We therefore represent images from the training set as normalized image vectors, from which the mean image is subtracted, \mathbf{x}_i ($i = 0, \dots, N-1$). These vectors form an image matrix $X \in \mathbb{R}^{n \times N}$

$$X = [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{N-1}], \quad (1)$$

where n is the number of pixels in the image and N is the number of training images.

The most straightforward way to calculate the eigenvectors is by solving the SVD of the covariance matrix $C \in \mathbb{R}^{n \times n}$ of these normalized vectors

$$C = XX^T$$


$$= [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{N-1}] \begin{bmatrix} \mathbf{x}_0^T \\ \mathbf{x}_1^T \\ \vdots \\ \mathbf{x}_{N-1}^T \end{bmatrix} = VLV^T. \quad (2)$$

The columns of V are the eigenvectors \mathbf{v}_i , $i = 0, \dots, N-1$, which form an orthogonal basis set. The values λ_i in the diagonal matrix L are the corresponding eigenvalues, which denote the overall variance that each of the eigenvectors encompasses.

Since the number of pixel elements n in an image is usually high, the computation of the matrix C is a time consuming task of high storage and computational demands. However, it is possible to formulate the equations in such a way that instead of decomposing C one can calculate the eigenvectors \mathbf{v}'_i , $i = 0, \dots, N-1$, of the inner product matrix $Q \in \mathbb{R}^{N \times N}$

$$Q = X^T X = \begin{bmatrix} \mathbf{x}_0^T \\ \mathbf{x}_1^T \\ \vdots \\ \mathbf{x}_{N-1}^T \end{bmatrix} [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{N-1}]. \quad (3)$$

Since the eigenvectors \mathbf{v}'_i are the solution of $X^T X \mathbf{v}'_i = \lambda'_i \mathbf{v}'_i$, we can calculate the eigenvectors of XX^T by $XX^T X \mathbf{v}'_i = \lambda'_i X \mathbf{v}'_i$ [4,19]. In this way, we



(a)



(b)

Fig. 3. (a) Panoramic image taken with a spherical mirror. (b) The same image warped to a cylindrical representation.




Fig. 4. Covariance of panoramic images in space. XY plane represents coordinates of the experimental environment. Measured and then interpolated is the covariance with image at $X = 440$ cm and $Y = 660$ cm.

obtain the eigenvectors \mathbf{v}_i of the covariance matrix C by projecting the \mathbf{v}'_i on the set of images,

$$\mathbf{v}_i = \frac{1}{\sqrt{\lambda'_i}} X \mathbf{v}'_i. \quad (4)$$

The eigenvectors corresponding to non-zero eigenvalues of the covariance matrix form the basis of the eigenspace, which reflects the training images in the subspace of a maximum of N dimensions. If we sort the eigenvectors according to the eigenvalues, we get an ordering which reflects the principal directions of variance in the learning set of images. We can therefore choose just a subset of p eigenvectors with the largest eigenvalues to be included in the model. Each image gets therefore optimally approximated in the least squares sense up to the degree of error we make with cutting off the least informative eigenvectors. Namely, every image \mathbf{x}_i in the training set projects into some point \mathbf{q}_i in the *eigenspace*, spanned by the selected eigenvectors [22], and can be approximated by

$$\mathbf{x}_i \approx \sum_{j=1}^p q_{ij} \mathbf{e}_j, \quad (5)$$

where p is the number of eigenvectors included in the eigenspace. The major advantage of the eigenspace method is that it allows to perform image matching in a low-dimensional space. Matching of images is usually performed by using some distance or similarity metric. In the image space, the city-block distance (L_1 norm) and the Euclidean distance (L_2 norm) are among the most common distance measures, while covariance and correlation are among the most common similarity measures.

In Fig. 4 we show a graph depicting the covariance with the panoramic image that was taken at the position $X = 440$ cm and $Y = 660$ cm. As we see, the covariance drops with the distance; position could therefore be inferred by finding the image in the training set that shows the largest covariance with the momentary panoramic view.¹ However, it can be shown that the covariance on original images is related to the L_2 norm on eigenspace projections of these images. It is therefore possible to match images in a much lower-dimensional space; usually it is done by

¹ Note that this is not necessarily true, as two distant locations may have very similar appearances. However, this can be solved by using the appearance model combined with a Markov localization procedure, as has been presented in [32].

representing all the images by their projections in the eigenspace and then by adopting the L_2 norm or the Mahalanobis distance between the projections as the distance metric. Furthermore, it is possible to densely interpolate the set of points in the eigenspace to obtain a hypersurface that represents an approximation of an arbitrary dense set of images [22].

To localize the robot we therefore have to find the appropriate parameters of the input image and then search for the nearest point on the hypersurface, which represents the model built on the basis of the images in the training set.

4. Eigenspace of spinning images

In the previous section we have described how a panoramic eigenspace model can be used for the localization task. As we pointed out, we perform recognition as finding the nearest projection of an image on the hypersurface. The measure of distance in our case is the Euclidean distance of corresponding projections in the eigenspace, which relates to the sum of squared differences between the two images in the image space. For the recognition algorithm to be successful it is therefore crucial that all of the images taken at

nearby positions are also neighbors in the eigenspace. However, for panoramic images this is only true when they are aligned in the same direction. In Fig. 4 we already showed a graph of correlation between equally oriented images, which depicts how correlation gradually decreases with the distance. But what happens if the images are taken in an arbitrary orientation? It is intuitive that the covariance of the neighbor images would drop. We can show with a simple experiment how the distribution of the covariance changes. In Fig. 5 we depict the covariance with the image taken at the position $X = 440$ cm and $Y = 660$ cm for the set where images were randomly oriented. One can see, that in this case covariance is not a reliable indicator of the distance between the locations.

Therefore, if we have access to an internal compass that indicates the orientation of the sensor, we can use its information in order to adequately orient the images. However, it is possible to overcome this problem by only using the visual input.

Several of the approaches that we reviewed in Section 2 dealt with the problem of in-plane rotations of panoramic images [2,9,24]. Driven by a desire to obtain a compact representation, the general tendency of the approaches applied until now has been to represent each position with just one image.




Fig. 5. Covariance of randomly oriented panoramic images in space. XY plane represents coordinates of the experimental environment. Measured and then interpolated is the covariance with image at $X = 440$ cm and $Y = 660$ cm.

Thus, different representations have been proposed that achieve some sort of rotational invariance, which results in a shift-invariance for cylindrical panoramic images. One way to achieve the invariance is by applying a transformation that produces the same output regardless of the shift in the input image. Such a transformation is autocorrelation (which in the particular case of cylindrical panoramic images can also be applied row-wise [2]). Another approach is to orient the images in a reference orientation, and thus preserve appearance [24].

We argue that while these approaches achieve shift-invariance, they cannot be used for robust localization. This is due to the fact that they use a global transform of the original images that is sensitive to local changes in the scene which are then spread across the transformed image.

What we propose in this paper is a representation that includes all the variations in rotation that we can simulate by shifting the cylindrical panoramic images row wise. Multiple orientations of the robot at the training positions are then represented as points in the eigenspace. As we show, the moderate increase in storage complexity pays off with robustness of position retrieval. Furthermore, the representation that we obtain by including all the rotated images in the training set is partially analytical; points corresponding to one position but under different orientations lie on analytical hypersurfaces, which makes localization easier. In this way we achieve rotation invariance without the need for searching for the correspondences or applying a global transform.

4.1. Eigenspace of a set of rotated versions of one panoramic image

Let us first assume that all of the images in the training set capture the same scene from a *single point of view*, but under different in-plane rotation. If we unwrap the panoramic image to a cylindrical projection, the rotation can be expressed simply as a row wise shift (Fig. 3). Since the images are uniformly rotated (shifted), each image can be generated by rotating (shifting) the original image \mathbf{x}_0 for $2\pi/N$.

Uenohara and Kanade [30] showed that in the case of the image set consisting of rotated examples of one original image, the inner product matrix Q is a symmetric Toeplitz matrix (Fig. 6). Since Q is also circulant,




Fig. 6. Visualization of the inner product matrix Q for rotated images captured at a single location.

lant, its eigenvectors \mathbf{v}'_i are not dependent on the contents of the images, while eigenvalues can be calculated without decomposing Q [30]. Q is of the form

$$Q = \begin{bmatrix} q_0 & q_1 & \cdots & q_{N-2} & q_{N-1} \\ q_{N-1} & q_0 & q_1 & \cdots & q_{N-2} \\ q_{N-2} & q_{N-1} & q_0 & q_1 & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ q_1 & \cdots & q_{N-2} & q_{N-1} & q_0 \end{bmatrix}. \quad (6)$$

It can be derived from the *shift theorem* [10], that the eigenvectors of a general circulant matrix are the N basis vectors from the Fourier matrix $F = [\mathbf{v}'_0, \mathbf{v}'_1, \dots, \mathbf{v}'_{N-1}]$

$$\mathbf{v}'_k = [1, w^k, w^{2k}, \dots, w^{(N-1)k}]^T, \quad k = 0, \dots, N-1, \quad (7)$$

where $w = e^{-2\pi j/N}$, $j = \sqrt{-1}$.

The eigenvalues can be calculated simply by retrieving the magnitude of the DFT of one row of Q

$$\lambda'_k = \sum_{i=0}^{N-1} q_i e^{-2\pi j(ik/N)}. \quad (8)$$

This interesting property also emphasizes the central point of the Fourier analysis, as it indicates that

the Fourier basis diagonalizes every periodic constant coefficient operator, in our case the circular shift operator [28].

Since $q_i = q_{N-i}$, our matrix is circulant symmetric, and therefore we can choose an appropriate set of real-valued orthogonal eigenvectors. As it turns out, the proper basis are the cosine functions from the real and the sine functions from the imaginary part of the Fourier matrix [28].

We can therefore compute the eigensystem of Q just by first computing the autocorrelation vector $[q_1, q_2, \dots, q_{N-1}]$, and then by calculating the λ'_i values, which should be afterwards sorted by decreasing magnitude. The eigenvectors \mathbf{v}'_i corresponding to k largest eigenvalues can then easily be selected from the corresponding basis vectors of the discrete cosine transform (DCT) [30]:

$$v'_{km} = \cos \left[\frac{\pi(2m+1)k}{2N} \right], \quad (9)$$

$$m = 0, \dots, N-1, \quad k = 0, \dots, N-1.$$

Thus, with the help of the DCT, it is possible to compute the basis vectors much more efficiently.

4.2. An extension to a set of rotated panoramic images acquired at different locations

When dealing with the problem of localization of a mobile robot [12], we need to encode P images from P locations, each of them being rotated N times (Fig. 2). In this case, we cannot directly apply the previous approach to the calculation of eigenvectors of circulant matrices, since the inner product matrix A ,

$$A = X^T X = \begin{bmatrix} Q_{11} & Q_{12} & \cdots & Q_{1P} \\ Q_{21} & Q_{22} & \cdots & Q_{2P} \\ \cdots & \cdots & \cdots & \cdots \\ Q_{P1} & Q_{P2} & \cdots & Q_{PP} \end{bmatrix} \quad (10)$$

is composed of several circulant blocks Q_{ij} , which are, in general, not symmetric. An illustration of A for five locations is shown in Fig. 7. However, as we will show, it is still possible to calculate the eigenvectors without performing the SVD decomposition of A .

We have to solve the eigenvalue problem

$$A\mathbf{w}' = \mu\mathbf{w}', \quad (11)$$




Fig. 7. Visualization of the inner product matrix A for five locations.

where (μ, \mathbf{w}') is the eigenpair of A . The fact that the matrix blocks Q_{ij} of A are circulant matrices is crucial. As it was already mentioned, every circulant matrix can be diagonalized in the same basis by Fourier matrix F . Consequently, all the matrices Q_{ij} have the same set of eigenvectors $\mathbf{v}'_k = [1, w^k, w^{2k}, \dots, w^{(N-1)k}]^T$, $k = 0, \dots, N-1$. We shall find the eigenvectors \mathbf{w}' of A among the vectors of the form

$$\mathbf{w}'_k = [\alpha_{k1}\mathbf{v}'_k, \alpha_{k2}\mathbf{v}'_k, \dots, \alpha_{kP}\mathbf{v}'_k]^T, \quad k = 1, \dots, N. \quad (12)$$

Eq. (11) can be rewritten blockwise as

$$\sum_{j=1}^P Q_{ij}(\alpha_{kj}\mathbf{v}'_k) = \mu\alpha_{ki}\mathbf{v}'_k, \quad i = 1, \dots, P. \quad (13)$$

Since \mathbf{v}'_k is an eigenvector of every Q_{ij} , the equations simplify to

$$\sum_{j=1}^P \alpha_{kj}\lambda_{ij}^k \mathbf{v}'_k = \mu\alpha_{ki}\mathbf{v}'_k, \quad i = 1, \dots, P, \quad (14)$$

where λ_{ij}^k is an eigenvalue of Q_{ij} corresponding to \mathbf{v}'_k . This implies a new eigenvalue problem




Fig. 8. First 56 real and imaginary components of eigenvectors calculated for a set of rotated images representing 63 locations.

$$\Lambda \alpha_k = \mu \alpha_k, \quad (15)$$

where

$$\Lambda = \begin{bmatrix} \lambda_{11}^k & \lambda_{12}^k & \cdots & \lambda_{1P}^k \\ \lambda_{21}^k & \lambda_{22}^k & \cdots & \lambda_{2P}^k \\ \cdots & \cdots & \cdots & \cdots \\ \lambda_{P1}^k & \lambda_{P2}^k & \cdots & \lambda_{PP}^k \end{bmatrix} \quad (16)$$

and

$$\alpha_k = [\alpha_{k1}, \alpha_{k2}, \dots, \alpha_{kP}]^T. \quad (17)$$

Since A is block-symmetric, Λ is symmetric and we have P linearly independent eigenvectors α_k which provide P linearly independent eigenvectors \mathbf{w}'_k in (12). Since the same procedure can be performed for every \mathbf{v}'_k , we can obtain $N \cdot P$ linearly independent eigenvectors of A .

It is therefore possible to solve the eigenproblem using N decompositions of order P . Since P is usually small in comparison to the total number of images $P \cdot N$, this method offers a similar improvement as the method in [30].

However, by looking at the properties of the circulant matrices one can deduce, that this method works only if we use the complex Fourier basis as the eigenvector set for the circulant matrix. In fact, this set of

basis vectors is the only common eigenspace for all the submatrices Q_{ij} from A .

Those complex eigenvectors \mathbf{w}'_k , that are not of zero frequency (gain), come in conjugated pairs. We can therefore reduce their number by half by taking only one eigenvector of a pair. In Fig. 8 are depicted the real and the imaginary parts of the final eigenvectors \mathbf{w}_i . It can be seen that the eigenvectors have a structure resembling row wise harmonic functions (pairs have the same frequency), however, every row may differ from the other in the amplitude and the phase of the signal. Indeed, these eigenimages carry the most of the pictorial information, since this set of basis functions is still the optimal linear approximation of the learning set.

To represent each training image in the eigenspace, we have to estimate its coefficients by projecting the image on the basis vectors. Since our eigenvectors are complex, the coefficient vectors are also complex, and can be viewed as points in the complex eigen-subspace. In Fig. 9 we can see the coefficients for an image and its rotated siblings. The rotation (shift) of the image results in the change of angle of the coefficient vector in the complex eigenspace, while the magnitudes of the coefficients do not change.

Because of this property, it is not necessary to store the coefficients for all the rotated images. In fact, if we




Fig. 9. Complex coefficients corresponding to the first 10 eigenvectors for one location of the training set. Depicted is the angle between the coefficients of two consecutive rotated images for the first eigenvector.

have just one representative coefficient for one viewpoint, all the others can be generated on the fly with a simple rotation in the complex plane.

Let \mathbf{x}_i be the i th image and \mathbf{x}_r , $r = i + 1, \dots, (i + N - 1)$, a set of its rotated siblings. Images are evenly rotated, each by $2\pi/N$ radians relative to its predecessor.

The coefficient q_{ij} of the image \mathbf{x}_i can be calculated as $\langle \mathbf{x}_i, \mathbf{e}_j \rangle$. Let the j th coefficient of the image $i + 1$ be $q_{(i+1)j}$. The angle between the two coefficients is

$$\Delta\phi_{ij} = \arctan \frac{\operatorname{Re}(q_{(i+1)j} - q_{ij})}{\operatorname{Im}(q_{(i+1)j} - q_{ij})}. \quad (18)$$

Every other coefficient q_{rj} can now be calculated as

$$q_{rj} = q_{ij} e^{2\pi\sqrt{-1}(r-i)\Delta\phi_{ij}}, \quad r = i + 1, \dots, (i + N - 1), \quad j = 1, \dots, K. \quad (19)$$

It is therefore easy to handle rotation using a set of complex eigenvectors, since the coefficients, also being complex themselves, can intuitively be handled in the recognition stage. However, if one wishes to obtain a real basis (this is possible since A is symmetric), it can be done in a straightforward procedure by using normalized vectors taken from the real and imaginary parts of complex eigenvectors. In this way, also the image coefficients are real, their values being the same as the real and imaginary parts of the complex counterparts. In Fig. 10 we depict a three-dimensional subspace, where the first three coefficients are used to illustrate the projections of images representing three locations. As one can see, coefficients for a single location, representing one panoramic image in several orientations, form a regular harmonic closed loop.




Fig. 10. Coefficients corresponding to the first three eigenvectors for locations 21, 29 and 31 of the training set in 25 orientations.




Fig. 11. Calculating the coefficients from a set of linear equations based on subset of points.

5. Making localization robust

To localize the robot we have to calculate the parameters of the momentary input image and then search for the nearest point in the eigenspace. Let us now for simplicity of notation assume that the coefficients are real numbers. The standard method to recover the parameters is to project the image vector \mathbf{y} onto the eigenspace [22]:

$$q_j(\mathbf{y}) = \langle \mathbf{y}, \mathbf{e}_j \rangle, \quad j = 1, \dots, p. \quad (20)$$

However, the way the parameters are calculated is non-robust and thus not accurate in the case when the input image locally deviates from the image approximated in the environment map. These deviations can sometimes be a side effect of the design of the panoramic sensor (e.g., self-occlusion of the camera holder). Also they may be caused by occlusions due to objects moving in the environment. Since panoramic images capture a wide field-of-view it is almost impossible to avoid any disturbances during the localization (operation) phase, thus it is essential that a representation of the environment map enables efficient robust matching.² For example, if we imagine a mobile robot roaming around with a model of a static environment, every change in the environment, such as displaced objects or people walking around can result in severe variations with respect to the original images learned at the training stage.

To overcome the erroneous calculation of image parameters when the visual content partially deviates from the learning examples, we propose to use the robust approach [16], that, instead of using the inner product (Eq. (5)), generates and evaluates

a set of hypotheses r as subsets of image points $r = (r_1, r_2, \dots, r_k)$. In fact, the coefficients can be retrieved by solving a set of linear equations on $k = n$ points:

$$x_{r_i} = \sum_{j=1}^n q_j(\mathbf{x}) e_{j r_i}, \quad 1 \leq i \leq n, \quad (21)$$

where x_{r_i} represents the point r_i of the image \mathbf{x} , and $e_{j r_i}$ represents the point r_i of the eigenvector \mathbf{e}_j , while $q_j(\mathbf{x})$ is the corresponding j th coefficient of \mathbf{x} . The principle of such computation is clearly illustrated in Fig. 11.

By selecting only p , $p \leq n$, eigenimages as our basis we have to solve an over-constrained system in a robust way, so that the solution set of parameters minimizes

$$E(\mathbf{r}) = \sum_{i=1}^k \left(x_{r_i} - \sum_{j=1}^p q_j(\mathbf{x}) e_{j r_i} \right)^2. \quad (22)$$

We solve the system on k , $k > p$, points, where k is significantly smaller than the total number of image points. The initial set of points is randomly initiated at first. After solving of the set of equations, we first perform an α -trimming step, in order to allow only the points on which the error is arbitrary small to contribute to the further computation of the parameters. The whole procedure then iterates either until the error gets small enough or until the number of remaining points meets some boundary condition. As we can see in Fig. 12, at this stage most of the points in the occluded regions are excluded from the computation.

To increase the probability of avoiding points that are noise or represent occlusion, several different subsets of points are generated, resulting in multiple hypotheses. A hypothesis consists of a set of parameters, an error vector ϵ calculated as the squared

² Robustness is defined as the extent of the ability of a method to give expected results despite the deviations of the input data.




Fig. 12. Crosses denote points that contribute to the generation of a hypothesis after three steps of α -trimming. Image is 60% occluded with vertical bars.

difference between the original image and its reconstruction $\sum_{j=1}^P q_{ij} \mathbf{e}_j$, and the domain of compatible points that satisfy an error margin constraint. These hypotheses are then subject to a selection procedure, based on the *Minimal Description Length* principle, as described in [16].

It is important to note at this point that the method is based on the assumption that images have not been preprocessed in a way that would prevent to correctly evaluate the outliers. As an example we can mention Aihara et al. [2] who used row-autocorrelated images, where error from outliers spreads globally over the transformed image. Further, the evaluation of the hypotheses is based on the reconstruction of the image, which gets evaluated against the original. Obviously, this also prohibits the use of global transforms.

6. Experimental results

In this section we give some experimental results evaluating the performance of robot localization in realistic environments. In Fig. 13 we can see an example of an office environment with marked positions where learning images were taken during the training phase (denoted with squares). We used two different mirrors in the catadioptric camera. The first mirror was of a spherical shape, while the second mirror had a hyperbolic reflective surface of an $r = 1.9$ ratio.³ The round images depicted in Fig. 3(a) were acquired with the spherical mirror setup, while the images in Fig. 14




Fig. 13. Map of $P = 62$ locations, where images were taken for training the robot. Images shown corresponding to locations 7, 19, 37 and 62, respectively.

were acquired using the hyperbolic mirror. During our experiments we did not notice any significant difference in the performance of the two mirror shapes.

In Table 1 we compared the times required for building the eigenspace by

- (1) using the standard decomposition of the correlation matrix XX^T ,
- (2) by calculating the decomposition of the inner product matrix $X^T X$, and
- (3) by using our approach.

The tests were made for images of dimensions 40×68 , the latter being the width of the image. Each image was therefore rotated 68 times, i.e., for 40 locations we got 2720 images. Since this is also the number of image elements, this is the border case when the covariance matrix is of the same size as the inner product matrix, and the complexity of the SVD method reaches its upper bound.

Table 1
Timings for building of eigenspace (in seconds)

Locations (P)	XX^T	$X^T X$	CPLX
10	2507.3	55.8	16.1
20	2569.6	429.2	105.3
30	2634.8	1400.3	312.4
40	3007.7	3252.3	853.2

³ For a review of catadioptric cameras see [8].




Fig. 14. Four examples of the robot's view from the path where 161 images were taken. Some of the images in the testing set contain people walking in front of the camera.

Once the eigenspace is built, we can interpolate the memorized coefficients in order to form a hypersurface that generalizes the acquired knowledge in order to represent the intermediate locations. An example of an interpolated hypersurface for the real coefficients corresponding to the first eigenvector is shown in Fig. 15 (coefficients are shown for one rotation only).

In the case of complex coefficients, if we memorized only one coefficient vector per location, we can search among the scores of N projected coefficients, i.e., we first search for the nearest neighbor of the coefficient vector corresponding to the input image and then we rotate this coefficient vector by Φ , where Φ is a vector of estimated $\Delta\phi_k$ values (Eq. (18)).

6.1. Using shift invariant properties of complex coefficients

The time complexity of the search can be quite high if the parameter space is large. It is, however, easy to decrease it dramatically by using a two-step search.

When using a complex eigenspace, a closer look at the coefficients reveals that the modulus of the coefficient vectors does not change for the coefficient of the original image and its rotated siblings. If \mathbf{q} is the coefficient vector of an input image, we first calculate its modulus vector $\mathbf{m} = [m_1, m_2, \dots, m_k]$

$$m_i = \sqrt{\text{Re}^2(q_i) + \text{Im}^2(q_i)}, \quad i = 1, \dots, k. \quad (23)$$




Fig. 15. Interpolated hypersurface for the coefficients of the first eigenvector. Coefficients belonging to the images from the training set are depicted as empty circles.




Fig. 16. Images from the training set.

In the first search we make a list of coefficients on the hyperplane which have a modulus vector close to \mathbf{m} . These points and their parameters are then the only candidates for the second search. As our experiments show, this strategy dramatically speeds up the algorithm.

The estimation of the momentary position can be done in a few seconds on a Pentium III computer running MATLAB, when the model of an office environments is represented with a 10 cm grid. This is the case when we search for the right coefficient globally, i.e., by exploring all the coefficient space. In a practical application, one could limit the searchable space by considering only the positions that are within reach from the previously estimated position by accounting the local odometry of the robot. A global search can therefore be applied only at the very beginning of the navigation.

6.2. Results of localization

In this section we will first present experiments on localization of a mobile robot by using a standard approach with images oriented in a reference orientation. The next experiment, where an eigenspace of spinning images will be used, will demonstrate the effectiveness of our approach. We will continue with experiments that demonstrate robustness in the presence of occlusions. Finally, we will demonstrate the behavior of the method in different types of environments.

First we present the results of localization of a mobile robot equipped with a panoramic camera in an environment of roughly $6\text{ m} \times 9\text{ m}$. The training images were taken on a $60\text{ cm} \times 60\text{ cm}$ grid of 62 locations, denoted by squares in Fig. 13. Thirty-five of them can be seen in Fig. 16. For the first experiment an eigenspace was built by taking only one image for every position.



Fig. 17. Images from the testing set.

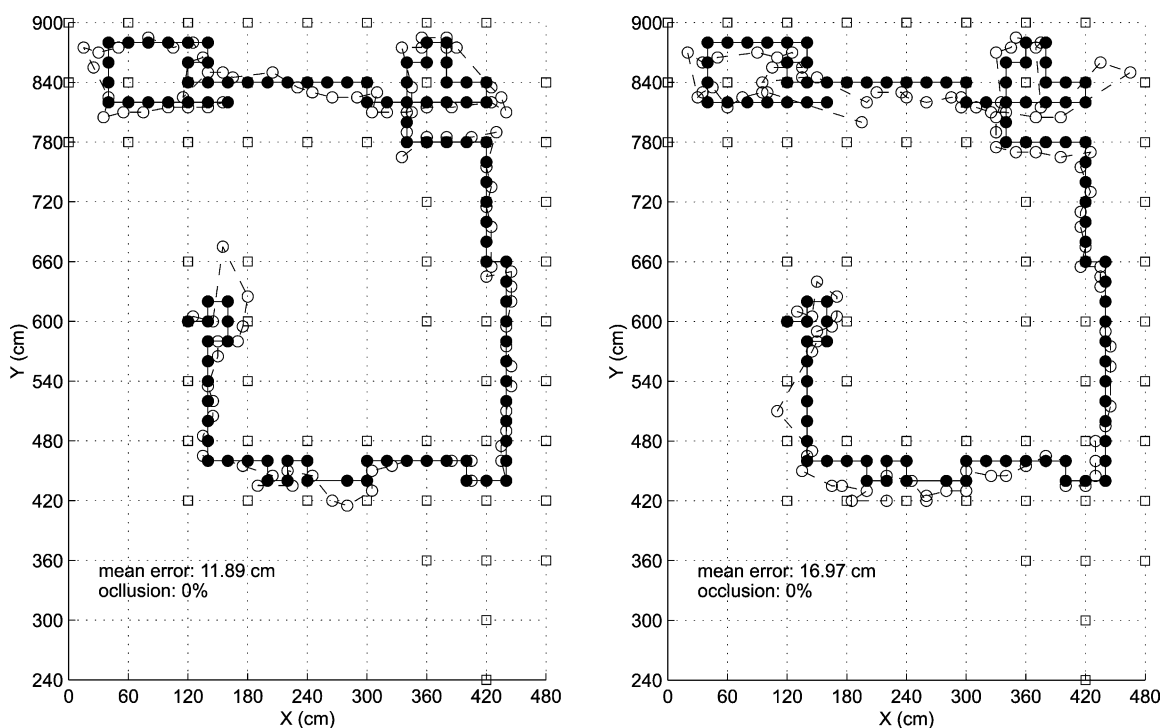


Fig. 18. Results of localization for the standard eigenspace approach using a 10-dimensional eigenspace (left) and the eigenspace of spinning images using 10 complex eigenvectors (right). Black dots denote the test locations, while empty dots denote the estimated locations of the robot.




Fig. 19. Mean error and standard deviation of localization for the standard and for the robust method.



Fig. 20. 60% occluded image. The occlusion is represented as vertical bars taken from an arbitrary gray-scale image.

In order to align the images as if the orientation of the robot was known, all of the images were virtually oriented in the same direction. Localization tests were then performed by taking an eigenspace of 10 dimensions onto which a set of 100 equally oriented test images was projected (Fig. 17). The position of the robot was then deduced from the calculated parameters for each image in the test set. In Fig. 18 (left), the testing positions are denoted by filled circles, while the recovered positions are denoted by blank circles. The average error of localization was 12 cm.

The next experiment was performed in the identical environment by using an eigenspace of spinning images. Each of the 62 images in the training set was

rotated 50 times, so that the interval of rotation was 7.2° . Tests were performed on a complex eigenspace of 10 dimensions and for the same testing set of 100 images (Fig. 17). To thoroughly test the rotational invariance, the images in the testing set were now rotated randomly in order to simulate multiple orientations of the sensor. The map with the results can be seen in Fig. 18 (right). Due to an increase of the training data, the dimension of the eigenspace needed to obtain comparable results increases (the effective dimension of the eigenspace with 10 complex eigenvectors is roughly 20, since every complex vector, except those whose imaginary component equals zero, has two components), however, as we can see from the map, by doubling the dimension we get a comparable average localization error of 16 cm.

6.3. Experimental results for the robust method

Since in the original set of images there is no significant occlusion (the only disturbances are some changes in the local illumination of the windows area), the standard method has a small error, mostly coming from the inaccuracy of the model itself. In fact, as it can be seen from the graphs in Fig. 18, the mean error of localization is between 11 cm and 16 cm for 0% occlusion.

The performance of the robust method at higher levels of occlusion noise is compared to that of the

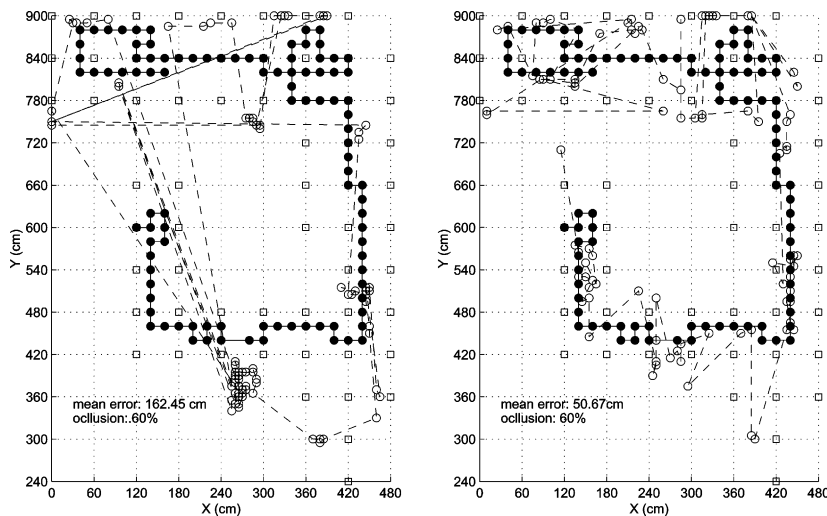


Fig. 21. Localization on an imaginary path of 100 images at 60% occlusion. Left: standard method; right: robust method.

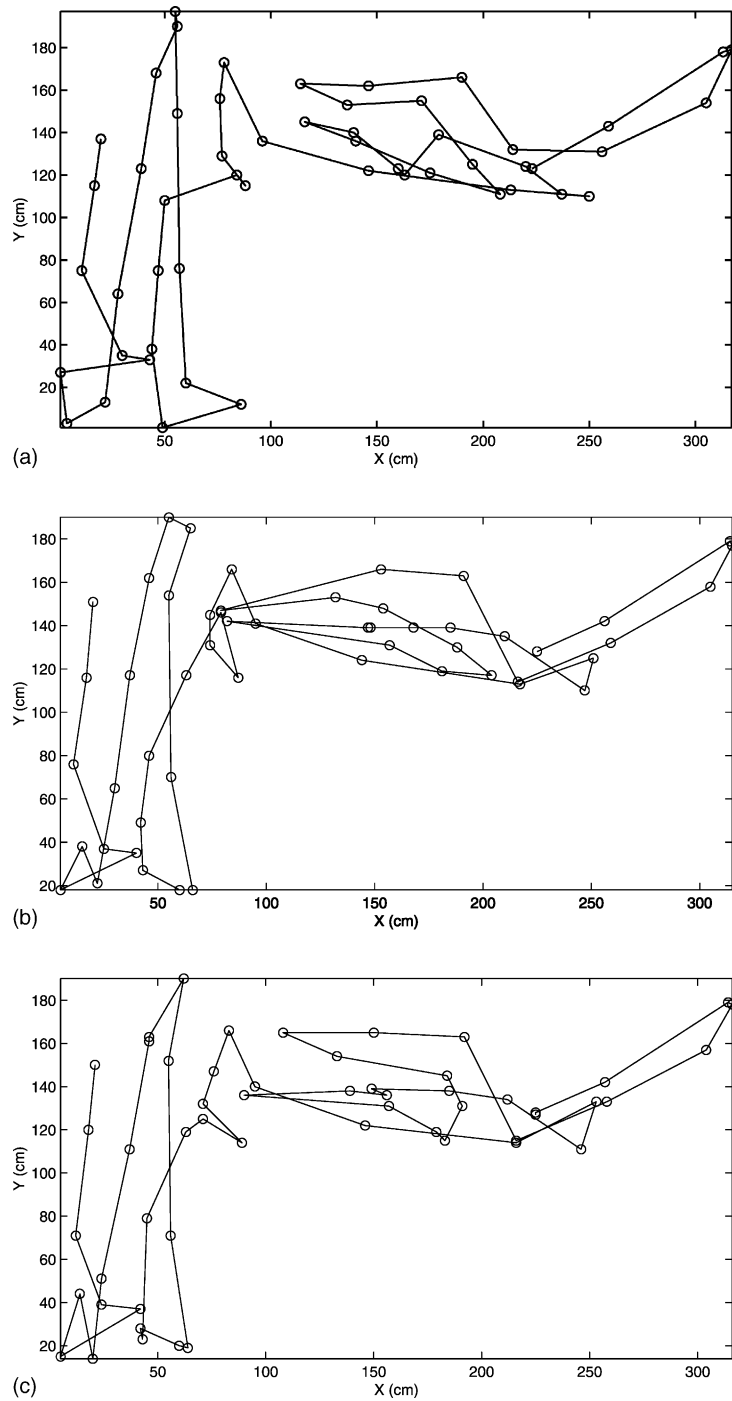


Fig. 22. Localization on an imaginary path of 54 images. (a) True positions, (b) positions estimated by the standard projection and (c) positions estimated with the robust algorithm.

standard method in Fig. 19. We can see a significant improvement in precision as a result of applying the robust method. Even in situations when more than half of the surrounding is occluded, as in Fig. 20, the robust method retrieves positions that are reasonably close to the correct ones. This can clearly be seen in Fig. 21. On the left we can see that the standard method breaks down because of the occlusion in the testing data, while the results of the robust estimator on the right show quite regular localization results with mean error under 60 cm, which is the resolution of the training set.

6.4. Results in different types of environments

To further test the algorithm we performed another set of experiments in a smaller area. The robot was roaming in a space of $2\text{ m} \times 3\text{ m}$ and taking images at predefined time intervals. This resulted in a set of 162 images that were then divided into two sets. The first set containing 108 images was used for training, while the remaining 54 images were used for testing the localization. In Fig. 22(a) we can see the measured locations of the testing images. The localization was first performed using a non-robust estimation of image parameters, which resulted in a set of estimated positions depicted in Fig. 22(b). The only disturbance in this set of images is the holder of the mirror, which represents a minor self-occlusion. Nonetheless, the re-

sults get considerably improved by using the robust method (Fig. 22(c)). In a smaller space the relative value of the error increases, however, it is still around 10 cm on the average. This performance is still satisfactory, as these are just the direct estimates, without using any of the knowledge on neighboring measurements to improve accuracy.

Next we questioned, how does the method perform when training images are taken at close intermediate intervals, so that the variance of the visual information captured in the training set is relatively small since panoramic images get very similar in appearance. In order to test this, we built a spinning images eigenspace by sampling locations that were no more than 15 cm apart. Then we sampled 40 test images, lying close together in a radius of approximately 1 m; these locations are denoted by empty circles in Fig. 23. The final recovered positions are denoted in the same figure as dots. The graph shows that even at a smaller scale the localization is stable and reliable.

Another experiment was made in a large entrance hall, which measures several dozens of meters in each direction. We acquired images on a path which measured approximately 6 m in length. Afterwards, we selected every fourth image as a training example. Finally, we run the localization algorithm on all the 161 images using a 10-dimensional subspace. The average error in this case was only 4 cm. Both the original and the recovered path can be seen in Fig. 24.




Fig. 23. Localization on a dense set of 40 images with the robust algorithm.




Fig. 24. Localization on a path of 161 images, where 41 images were used for training. (a) Original path and (b) recovered path.

7. Conclusion

In this paper we presented a framework for visual self-localization of mobile robots using an appearance-based parametric model build from panoramic snapshots of the environment. We proposed the “*eigenspace of spinning-images*” as a viable representation that solves many of the problems of the standard appearance-based techniques. It efficiently encodes the varying orientations of the sensor and the resulting eigenvectors provide an analytical framework for the computation of parameters where rotation is being handled easily. This allows for robust localization under arbitrary orientations of the robot in the case when no external source of information on orientation is available to align the images. We showed how to efficiently calculate the optimal subspace of a set of training snapshots even without actually decomposing the covariance matrix. We also introduced a robust method for the retrieval of parameters which enables robust localization even in the presence of noise and occlusion. In this way, the robot is capable of localizing itself in dynamic and unstructured environments.

Our current ongoing work is concentrated mainly on the incremental learning, which would enable us to bridge the gap between the learning and the training stage [3]. This coincides with many of the efforts in the

robotics community intended to provide frameworks that allow for simultaneous localization and modeling (SLAM).

Among the many advantages that an incremental algorithm brings we concentrate on the possibility of the dynamical building of multiple eigenspaces, where each of the eigenspaces would represent a locally coherent part of the environment [17]. Another problem of the classical eigenspace approach that we are trying to solve is the non-robustness of the learning phase. We are currently in the phase of developing algorithms for the robust building of the eigenspaces.

Lately, we achieved significant progress regarding the problem of changing illumination. As we showed, such a variation in appearance cannot be solved satisfactorily solely by the robust method of estimation of parameters. We therefore successfully applied a novel method that is based on filtering the eigenspace with a set of gradient filters [5].

The primary intent of the experiments presented is to show the capability of the appearance-based localization, showing results obtained solely by inferring from the momentary input image. In a real-world implementation, this system would preferably be integrated in a probabilistic framework which would account also for temporal coherence of the results, by using additional knowledge on the local robot movements. We therefore intend to put efforts in the

development of such a system as we believe that this will influence the stability and the precision of localization in a positive way.

Acknowledgements

The authors acknowledge the support from the Ministry of Education, Science, and Sport of the Republic of Slovenia (SI-A, SI-CZ, SI-USA, research program Computer Vision 506), the EU CEC (IST-2000-29375) CogVis project and the ECVision, EU (IST-2002-35454) research network.

References

- [1] Proceedings of the IEEE Computer Society Workshop on Models versus Exemplars in Computer Vision, Kauai, Hawaii, IEEE Computer Society, December 2001.
- [2] H. Aihara, N. Iwasa, N. Yokoya, H. Takemura, Memory-based self-localisation using omnidirectional images, in: Proceedings of the 14th International Conference on Pattern Recognition, IEEE Computer Society Press, Silver Spring, MD, August 1998, pp. 297–299.
- [3] M. Artač, M. Jogan, A. Leonardis, Mobile robot localization using an incremental eigenspace model, in: Proceedings of ICRA 2002, Washington, DC, IEEE Robotics and Automation Society, 2002, pp. 1025–1030.
- [4] D. Ballard, *Natural Computation*, MIT Press, Cambridge, MA, 1996.
- [5] H. Bischof, H. Wildenauer, A. Leonardis, Illumination insensitive eigenspaces, in: Proceedings of the International Conference on Computer Vision (ICCV'01), Vancouver, BC, vol. 1, IEEE Computer Society, 2001, 233–238.
- [6] M. Dill, R. Wolf, M. Heisenberg, Visual pattern recognition in *Drosophila* involves retinotopic matching, *Nature* 365 (1993) 751–753.
- [7] J. Gaspar, N. Winters, J. Santos-Victor, Vision-based navigation and environmental representations with an omnidirectional camera, *IEEE Transactions on Robotics and Automation* 16 (6) (2000) 890–898.
- [8] C. Geyer, K. Daniilidis, Catadioptric projective geometry, *International Journal of Computer Vision* 45 (3) (2001) 223–243.
- [9] H. Ishiguro, S. Tsuji, Image-based memory of environment, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Osaka, Japan, 1996, pp. 634–639.
- [10] B. Jähne, H. Haußecker, P. Geißler (Eds.), *Handbook of Computer Vision and Applications*, Academic Press, New York, 1999.
- [11] M. Jogan, A. Leonardis, Panoramic eigenimages for spatial localisation, in: Proceedings of the Eighth International Conference on Computer Analysis of Images and Patterns, Springer, Berlin, 1999, pp. 558–567.
- [12] M. Jogan, A. Leonardis, Robust localization using eigenspace of spinning-images, in: Proceedings of the IEEE Workshop on Omnidirectional Vision, IEEE Computer Society, June 2000, pp. 37–44.
- [13] M. Jogan, A. Leonardis, Robust localization using panoramic view-based recognition, in: Proceedings of the 15th International Conference on Pattern Recognition, Barcelona, Spain, vol. 4, IEEE Computer Society, September 2000, pp. 136–139.
- [14] S.P.D. Judd, T.S. Collett, Multiple stored views and landmark guidance in ants, *Nature* 392 (1998) 710–714.
- [15] B.J.A. Kröse, N. Vlassis, R. Bunschoten, Y. Motomura, A probabilistic model for appearance-based robot localization, *Image and Vision Computing* 6 (19) (2001) 381–391.
- [16] A. Leonardis, H. Bischof, Robust recognition using eigenimages, *Computer Vision and Image Understanding—Special Issue on Robust Statistical Techniques in Image Understanding* 78 (1) (2000) 99–118.
- [17] A. Leonardis, H. Bischof, J. Maver, Multiple eigenspaces, *Pattern Recognition* 35 (11) (2002) 2613–2627.
- [18] S. Maeda, Y. Kuno, Y. Shirai, Active navigation vision based on eigenspace analysis, in: Proceedings of the International Conference on Intelligent Robots and Systems (IROS'97), Grenoble, France, IEEE Computer Society, 2001, pp. 1018–1023.
- [19] H. Murase, S.K. Nayar, Visual learning and recognition of 3D objects from appearance, *International Journal on Computer Vision* 14 (1) (1995) 5–24.
- [20] H. Murase, S.K. Nayar, Detection of 3D objects in cluttered scenes using hierarchical eigenspace, *Pattern Recognition Letters* 18 (4) (1997) 375–384.
- [21] D. Murray, C. Jennings, Stereo vision based mapping and navigation for mobile robots, in: Proceedings of the IEEE International Conference on Robotics and Automation ICRA'97, Albuquerque, NM, April 1998, pp. 1694–1699.
- [22] S.K. Nayar, S.A. Nene, H. Murase, Subspace methods for robot vision, *IEEE Transactions on Robotics and Automation* 12 (5) (1996) 750–758.
- [23] S.K. Nayar, T. Poggio (Eds.), *Early Visual Learning*, Oxford University Press, Oxford, 1996.
- [24] T. Pajdla, V. Hlaváč, Zero phase representation of panoramic images for image based localization, in: Proceedings of the Eighth International Conference on Computer Analysis of Images and Patterns, Springer, Berlin, 1999, pp. 550–557.
- [25] L. Paletta, S. Frintrop, J. Hertzberg, Robust localization using context in omnidirectional imaging, in: Proceedings of the International Conference on Robotics and Automation (ICRA 2001), Seoul, Korea, 2001, pp. 2072–2077.
- [26] J. Ponce, A. Zisserman, M. Hebert (Eds.), *Object Representation in Computer Vision II*, Proceedings of the ECCV'96 International Workshop, Cambridge, vol. 1114, Lecture Notes in Computer Science, Springer, Berlin, 1996.
- [27] S. Se, D. Lowe, J. Little, Vision-based mobile robot localization and mapping using scale-invariant features, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'01), Seoul, Korea, May 2001, pp. 2051–2058.

- [28] G. Strang, The discrete cosine transform, *SIAM Review* 41 (1) (1999) 135–147.
- [29] M.J. Tarr, H.H. Bülthoff, *Object Recognition in Man, Monkey, and Machine*, Cognition Special Issues, MIT Press, Cambridge, MA, 1999.
- [30] M. Uenohara, T. Kanade, Optimal approximation of uniformly rotated images: relationships between Karhunen–Loeve expansion and discrete cosine transform, *IEEE Transactions on Image Processing* 7 (1) (1998) 116–119.
- [31] N. Vlassis, Y. Motomura, B. Kröse, Supervised linear feature extraction for mobile robot localization, in: *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000, pp. 2979–2984.
- [32] N. Vlassis, B. Terwijn, B. Kröse, Auxiliary particle filter robot localization from high-dimensional sensor observations, in: *Proceedings of the IEEE International Conference on Robotics and Automation*, Washington, DC, May 2002, pp. 7–12.



Matjaž Jogan received his B.Sc. and M.Sc. degrees in Computer Science from the University of Ljubljana in 1999 and 2002, respectively. He is currently a teaching assistant and a Ph.D. candidate at the Faculty of Computer and Information Science, University of Ljubljana. As a member of the Cognitive Vision Group at the Computer Vision Laboratory, he is active as a researcher in the field of cognitive vision. Within this area, he is particularly interested in subspace methods for visual learning, omnidirectional vision, and mobile robot localization.



Aleš Leonardis received his Ph.D. degree in Computer Science from the University of Ljubljana in 1993. Currently, he is an associate professor at the Faculty of Computer and Information Science at the University of Ljubljana. From 1988–1991 he was a visiting researcher in the General Robotics and Active Sensory Perception Laboratory at the University of Pennsylvania. Between 1995 and 1997 he was a postdoctoral associate at Vienna University of Technology, Austria. He was also a visiting professor at the ETH, Switzerland, and at the Technische Fakultät der Friedrich-Alexander-Universität in Erlangen, Germany. He is author or co-author of more than 120 papers published in journals and conferences, and he co-authored the book *Segmentation and Recovery of Superquadrics* (Kluwer, 2000). His research interests include robust and adaptive methods for computer vision, object and scene recognition, learning, and 3D object modeling. He is an associate editor of *Pattern Recognition*. Also, he is currently serving as the Slovenian representative in the Governing Board of the International Association for Pattern Recognition.