Information Sampling for Optimal Image Data Selection

Niall Winters

Computer Vision and Robotics Group,
Department of Computer Science,
University of Dublin, Trinity College,
Dublin 2 - Ireland.
Niall.Winters@cs.tcd.ie

José Santos-Victor

Instituto de Sistemas e Robótica,
Instituto Superior Técnico,
Av. Rovisco Pais, 1,
1049-001 Lisboa - Portugal.
jasv@isr.ist.utl.pt

Abstract. This paper presents a statistical method termed *Information Sampling* for selecting the most relevant data from an a priori set of images. These data could be a single pixel or a number scattered throughout an image. The main problem addressed is how to determine which image data points contain the most relevant information. As distinct from other techniques, we utilize the inherent information contained within the image set. We show that the most relevant data points, i.e. the most discriminatory, are those which allow reconstruction of a given image with the smallest amount of error.

We apply *Information Sampling*, yielding the most informative data and subsequently rank this data from most to least discriminatory. We show how *Information Sampling* can be applied to determine the qualitative position of a mobile robot in an indoor environment, using only the highest ranking data.

1 Introduction

In this paper we consider an enhancement to the current approaches for visually guiding a mobile robot through an indoor environment. We term our method *Information Sampling*. In our previous work [4, 16, 17], we used entire omnidirectional images to visually navigate the mobile robot to its destination. The qualitative position of the robot was calculated by matching the current image to an a priori acquired set, i.e. an appearance based approach to the problem. Appearance based methods have not only been applied to the task of robot navigation but to many areas of computer vision. They have been shown to be successful for a number of recognition tasks, including character recognition [7], tracking [1] and face recognition [14, 15].

Often, including our case, in order to compress large amounts of data, appearance based systems are built using Principal Component Analysis. Construction of such a system involves computing the eigenvectors (sometimes called eigenimages) of an a priori set of images. The variance of this set is captured by its first few eigenvectors. This low dimensional subspace [8], also known as an eigenspace, forms an orthonormal basis into which each image from the a priori set is projected. Once this eigenspace has been built, real time recognition of an unknown image is achieved by projecting it into the eigenspace and using a simple distance measure to find its closest match to the previously projected points.

1.1 Existing Methods

It should be noted that the method just described represents a “global approach” to the recognition problem in that the entire image is used for projection. In terms of object recognition using standard images, several authors have noted problems with this approach including sensitivity to occlusion, scale change and illumination. The problem of dealing with partial occlusion (in a bin-picking task) was investigated by Ohba and Ikeuchi [9]. Instead of projecting the entire image, as is usual, they proposed dividing each image into a number of smaller windows which they termed *eigenwindows*. Eigenspace analysis was then applied to each window. Their basic idea was that even if a number of the windows were occluded, the remaining ones would contain enough information to perform the given bin-picking task. As they point out, a very large number of image windows need to be stored in order for the method to obtain good results. For example, if one had an a priori set of 1000 images of size 256×256 pixels, and each window was 8×8 pixels in size, then one would require 1,024 windows to represent an image or 1,024,000 to represent the entire a priori set. Clearly the chances of one window, acquired at runtime being matched to a number of images from the a priori set is high. This could be due, for example to having many ambiguous regions within an image. As noted by Colin de Verdière and Crowley [2] this leads to the problem of deciding which eigenwindows contain decriminative information and therefore should be used in the recognition task. It is highly desirable that only the most *effective* windows are selected from each acquired image, and that only these chosen windows be matched to the a priori set.

As a solution to this problem, Ohba and Ikeuchi propose using three criteria to eliminate the redundant windows, namely: detectability, uniqueness and reliability. Colin de Verdière and Crowley reformulate the problem as a question of whether to used the set of eigenwindows selected by a particular interest operator or to use those windows selected from a predefined grid.

Interest points are local features where the signal changes two-dimensionally. A large number of interest point detectors exist [3], but the one that has shown the most repeatable results is the Harris detector [5]. Schmid and Mohr [13] used such a detector to determine where to compute local grayvalue invariants in addressing the problem of image retrieval from a large database. In our case, a similar approach to determining which windows are the most effective, would be to only use windows that contain a number of interest points above a certain threshold. This is the approach taken by Jugessur and Dudek [6]. The interest point detector they use is a symmetry based context free attention one from [10]. Unfortunately, for such an approach to be effective, the images are required to be highly textured. Our a priori images do not exhibit such a property, as they were acquired along a corridor for the purpose of visually guiding a mobile robot. Here, the environment consists of plain white walls and brown doors. Additionally, with interest operators stability is hard to guarantee when changes in illumination occur.

As an alternative to using interest operators to determine which eigenwindows to choose, one may use windows chosen from a pre-defined grid [2]. The first stage of this approach involves projecting *all* of the eigenwindows into the eigenspace. Since an image will contain a number of windows, it is represented in the eigenspace as a surface and a set of images are represented by a set of surfaces. Naturally, on projection of an eigenwindow many matches will occur. Thus, suppression of redundant windows is required. This is usually achieved by noting that a search for the closest point in the eigenspace produces too many matches. Alternatively, suppression can occur at the training stage given that a redundant window will be projected many times. Nevertheless, this approach still requires enough space and computational power to store and search for all of the eigenwindows.

1.2 Our Approach

Our approach to this problem is somewhat different from those outlined above and is based on a method by Rendas and Perrone [11]. We noticed that the above approaches do not make use of the inherent information available from the a priori set of images. In the case of [9, 2] all the eigenwindows from entire images are first collected, from which the relevant data must be found. In the case of [13, 6] their approaches require extraction of features from images *before* proceeding.

In contrast, our approach termed *Information Sampling* selects the most relevant data from a set of images without using eigenspace analysis or needing to apply interest operators. Theoretically, it can be applied on a *pixel-by-pixel basis to any type of image*, as outlined in Section 2.1. In this paper, for computational reasons, we use windows instead of pixels, extracted from omnidirectional images. Essentially, we can reconstruct an image using only the data selected by *Information Sampling* and then minimize the error associated with this reconstruction. We can rank the selected data and choose how much of it we wish to utilize for matching. *Information Sampling* selects quality image data as distinct from relying upon a large quantity of data. It is only after the ranking stage that we employ a local eigenspace approach to perform the qualitative recognition of a robot's position within an indoor environment.

The outline of this paper is as follows: Section 2 presents the *Information Sampling* method and underlying statistics. In Section 3, we show how to use a local eigenspace approach for robot navigation while in Section 4, we present our experimental results and discuss the complexity of our method. Finally, in Section 5 we conclude and give the future directions of our research.

2 The Information Sampling Method

As previously noted, our approach requires the use of a priori image data. We wish to make it clear that our method is *independent* of image type. For the experiments outlined in this paper images acquired from an omnidirectional camera with a spherical mirror, built in-house at the Instituto de Sistemas e Robótica, Lisboa were used. This camera was mounted on a Labmate mobile platform and images were captured as it traversed through a corridor environment. The system is shown in Figure 1. Once the images were captured, we determined which regions contained the most relevant information, i.e. which were the most discriminatory by applying *Information Sampling.* As a first step in explaining this process, Section 2.1 outlines the procedure for reconstructing an image, given only a small amount of data.

Fig. 1. Left: the omnidirectional camera. Right: the camera mounted on the mobile robot.

2.1 Image Reconstruction

We assume that the images captured by the robot's camera can be modeled as a random vector I, characterized by a Gaussian distribution with mean $\bar{I}$ and covariance Σ_I:

$$I \sim \mathcal{N}(\bar{I}, \Sigma_I) = p(I)$$

Usually, one can take an ensemble of images of the environment $[I_1 \dots I_m]$, which can be utilized for computing $\bar{I}$ and Σ_I, so that $p(I)$ can be computed *a priori.* When the robot is navigating, we assume that the observations, d, consist of a selection of (noisy) image pixels (or sub-regions), rather than the entire image. Accordingly, the observation model can be expressed as:

$$d = SI + \eta \tag{1}$$

where d stands for the observed data and the measurement noise, η is assumed to follow a Gaussian distribution with zero mean and covariance, Σ_n. We further assume that I and η are independent. The selection matrix, S is composed of a series of ones and zeros, the ones corresponding to the data points extracted from an image. We select a number of pixels to test by moving the set of ones in the selection matrix.

Having prior knowledge of I, in the form of a statistical distribution, $p(I)$, the problem now consists of estimating the (entire) image based on partial (noisy) observations of a few pixels, d. This problem can be formulated as a *Maximum a Posteriori* estimation of I. The posterior probability can be determined from Bayes rule as follows:

$$p(I|d) = \frac{p(d|I)p(I)}{p(d)} \tag{2}$$

where $p(d|I)$ is the likelihood of a pixel (or set of pixels) given a known image, I; the prior distribution is denoted by $p(I)$ and is assumed to have been learnt *a priori.* With this information we calculate the maximum a posteriori estimate of an image, $\hat{I}_{MAP}$ [12] as follows:

$$\hat{I}_{MAP} = \arg\max_I p(I|d) = (\Sigma_I^{-1} + S^T \Sigma_n^{-1} S)^{-1} (\Sigma_I^{-1} \bar{I} + S^T \Sigma_n^{-1} d) \tag{3}$$

Thus, $\hat{I}_{MAP}$ is the reconstructed image obtained using the pixel (or set of pixels), d. Notice that by combining the prior image distribution with the statistical observation model, we can estimate the entire image based on the observation of a limited number of pixels.

2.2 Choosing the Best Data: Information Windows

Once we have reconstructed an image using the selected data, we can compute the error associated with this reconstruction. The error covariance matrix, Σ_{error} is given by:

$$\Sigma_{error} = \mathrm{Cov}(I - \hat{I}_{MAP}) = (\Sigma_I^{-1} + S^T \Sigma_n^{-1} S)^{-1} \tag{4}$$

Of course, the quality of the estimate, and the "size" of Σ_{error} depend not only on the observation noise, η but also on the observed image pixels, as described by the selection matrix, S. Equation (4) quantifies the quality of an estimate obtained from using a particular set of image pixels. In theory, we can evaluate the *information content* of any individual image pixel or combination of pixels, simply by selecting an appropriate selection matrix, S, and determining the associated Σ_{error}.

This problem could be formulated as an experiment design process [12] , in which we look for the optimal selection matrix S^* that minimizes (in some sense) the error covariance matrix. If we take the determinant of Σ_{error} as an indication of the "size" of the error, the optimal selection of image pixels would be given by:

$$S^* = \arg\min_S \{ \, det((\Sigma_I^{-1} + S^T \Sigma_n^{-1} S)^{-1}) \, \} \tag{5}$$

In practice, to avoid computing the inverse we define the following equivalent optimization problem in terms of a modified uncertainty metric, U:

$$U = -\log\{ \, det(\Sigma_I^{-1} + S^T \Sigma_n^{-1} S) \, \} \, ; \quad S^* = \arg\min_S U \tag{6}$$

So far, we have described *Information Sampling* as a process for (i) reconstructing an entire image from the observation of a few (noisy) pixels and (ii) determining the *most relevant* image pixels, S^*, in the sense that they convey the most information about the image set.

Unfortunately, determining S^* is computationally impractical since we would have to compute Σ_{error} for all possible combinations of pixels scattered throughout the image. Instead, we partition the image into non-overlapping square windows of $(l \times l)$ pixels. We term these regions *Information Windows*, denoted by $\mathbf{w} = [w_1 \dots w_n]$.

By using equation (6), we can rank *Information Windows* or combinations of such windows, in terms of their information content. Again, as searching for all possible combinations of windows within the image, in order to minimize equation (6), would be computationally intensive, we instead use two sub-optimal (greedy) algorithms. These algorithms are described in Section 4.1.

Notice that the information criterion is based on the entire set of images and not, as with other methods, on an image-by-image basis. For instance, a highly textured image region would only be selected if it varied significantly from one image to the next.

3 Local Appearance Space

While visually guiding our mobile robot could conceivable be achieved by directly matching the reconstructed image, $\hat{I}_{MAP}$ to the set of omnidirectional images, this would be computationally expensive. Instead, by compressing data using Principal Component Analysis (PCA), robot position can be determined in real time. When using PCA, it is usual for entire omnidirectional images to be utilized for both eigenspace construction and image projection, i.e. a global approach (see our previous work [4, 16, 17]. Figure 2 shows an image acquired by the robot on a run down a corridor and the reconstruction of the closest image from the a priori set. As one can see, the correct qualitative position of the robot within its environment was calculated.

We wish to improve upon this method by projecting *only* the most relevant information obtained by *Information Sampling* into a local appearance space. In this way we significantly reduce the number of projected windows, thus immediately reducing the level of possible ambiguity. Additionally we reduce even further the amount of data used for matching. The local appearance space has an orthonormal basis of eigenvectors of size $(l^2 \times 1)$, where l is the length of the side of a square information window.

Following the standard PCA approach, and using Singular Value Decomposition (SVD) we can determine the eigenvectors (sometimes called eigenwindows), $\mathbf{e}_j$, and eigenvalues, λ_j, of the covariance matrix Σ_{local} of the windows selected from the set of omnidirectional images, $\mathbf{I}^w$. We denote the selected windows from each image, $\mathbf{w} = [w_1 \dots w_m]$. We denote α_j as the vector of co-efficients obtained by projecting each window

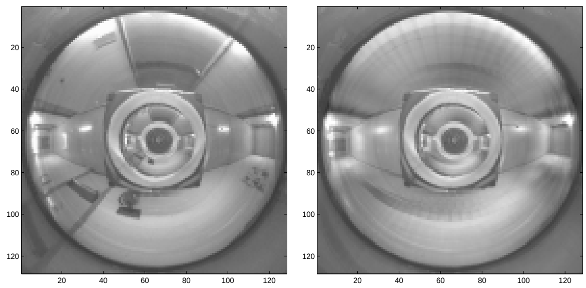

Fig. 2. The input (left) and retrieved (right) omnidirectional images are very similar.

from $\mathbf{I}^w$ into the local eigenspace. We can reconstruct an entire unknown *image*, $\hat{A}_{MAP}$ by replacing d in equation (3) by w_k and reconstruct its associated *window*, $\hat{g}$ from PCA as follows:

$$\hat{g} = [e_1 \dots e_q]\alpha_j \tag{7}$$

4 Experimental Results

Information Sampling has been tested on omnidirectional images acquired in a simple indoor office environment. Processing was carried out off-line on a Celeron 333MHz PC using Matlab. The a priori set of ninety omnidirectional images were obtained every 50 cm and ordered according to the direction of motion of the Labmate mobile platform. Each image was acquired at a resolution of 768×576 pixels using a Tekram acquition board. Once each image was acquired it was filtered and subsampled to an image resolution of 32×32 pixels. Additionally, the images were rewarped to a panoramic viewpoint, filtered and subsampled to 8×80 pixels in size. This representation explicitly takes advantage of the polar structure of the images. The reason for such small image sizes relates to the complexity of determining the error covariance matrix, Σ_{error} in equation (4). For example, if one wished to use images of 128×128 pixels, the computation of Σ_{error} would require the calculation of a matrix of $16,384 \times 16,384$ elements in size, a prohibitively large computation.

Each information window was chosen to be 8×8 pixels in size, thus giving 16 non-overlapping information windows per omnidirectional image. For panoramic images, the windows were 8×10 pixels in size, thus giving ten non-overlapping information windows per image. For image reconstruction using equation (3), these windows were of adequate size. As described in Section 3, we use a local appearance space to determine the qualitative position of a mobile robot within its environment. Additional positioning experiments were undertaken using eigenwindows of 32×32 pixels in size, extracted from omnidirectional images of 128×128 pixels in size. In order to locate the 32×32 pixel sized information windows, we used the ratio: 64:1,024 = 1,024:16,384.

4.1 Ranking the Information Windows

We tested two algorithms to rank the information windows: Combinatorial Search and Simple Search.

1. **Combinatorial Search** We first search for the best *Information Window*. Then, the search for the next best window is made keeping the first window fixed, thus locating the best *pair* of windows. As the method continues it determines the best triplet of windows, etc. If we denote n as the number of windows within an image, this method requires the evaluation of equation (6), $n!$ times. The method automatically groups the information window(s) into a single window, a pair of windows, a triplet of windows etc.
2. **Simple Search** This is a faster search algorithm. We rank each of the information windows *independently*. In this case, equation (6) has only to be evaluated n times. As distinct from the combinatorial method, if we wish to group the best (single, pair, triplet etc. of) windows we must do it manually based on the initial ranking.

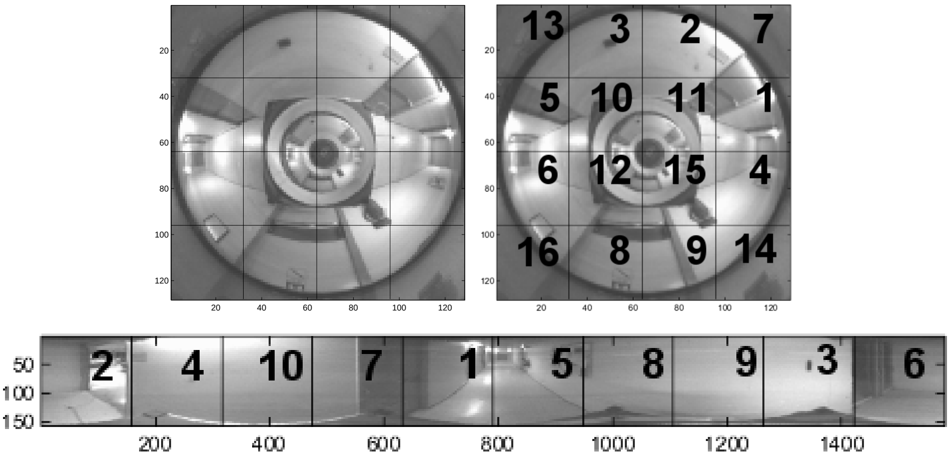

Fig. 3. Top Left: The 16 non-overlapping information windows. Top Right: Those windows ranked according to the amount of information they contain using Simple Search. Bottom: Information window ranking when using panoramic images.

These methods were chosen given the complexity of searching for the optimal solution. Figure 3 (top left) shows the information windows available for selection and (top right) these information windows, *individually ranked* from the most (number 1) to the least discriminating (number 16) using Simple Search. Figure 3 (bottom) shows the ranking when using panoramic images.

To give an intuitive idea of the *Information Sampling* method, we present the following example. In any of the omnidirectional images in this paper, the robot is in the centre of each image. Any information window which contains the robot is not a discriminating one and so it follows that such a window should have a lower ranking. As shown in Figure 3 (top right), this proves to be the case: the four information windows which contain the robot are ranked from numbers ten to fifteen. Additionally, the four windows at the periphery of the image also have a low ranking, since they only contain a portion of the omnidirectional image. It should be noted that the corridor in which the a priori set of images were acquired has a number of offices on one side (the top half of the omnidirectional images) and only a single door and notice-board on the other (the bottom half of the omnidirectional images). Thus, as the robot travels down the corridor more information change occurs in the top half of the omnidirectional images. Again, this is borne out by the window ranking, where the three highest ranking information windows are all in the top half of the omnidirectional image.

Figure 4 shows the graphs of the information windows, obtained from omnidirectional images, ranked using Simple Search (left) and Combinatorial Search (right). In both cases, the x-axis corresponds to the window ranking, from first to sixteenth and the y-axis corresponds to the uncertainty metric, U calculated using equation (6). The numbers along the graph line correspond to the 16 non-overlapping information windows per omnidirectional image. For example, using Simple Search the left graph tells us that the eighth information window exhibits the lowest uncertainty value and so is individually ranked in first position, while the third window, having a higher uncertainty value, is individually ranked in second position etc.

Using Combinatorial Search the right graph tells us that the eighth window is ranked in first position. This window is then fixed and the best *pair* of windows, in this case the eight plus the third, are found. Thus, the third window contains the next best amount of information and is ranked in second position. Using Combinatorial Search the next best window added at each stage matches the window rank chosen by Simple Search.

Combinatorial Search continues until all windows have been combined. As can be see from Figure 4 (right) each *combination of information windows* exhibits a lower uncertainty measure than the previous one. Intuitively, this makes sense as the more information available, the better the image reconstruction should be. However, the payoff for using many information windows is not significant as can be seen from the small drop in uncertainty. This result is also borne out by Figure 5, as detailed in Section 4.2. Clearly, the fact that the highest ranking information window is not only the most relevant, but is the most relevant by a significant factor, is the reason why we need use only it for reconstruction.

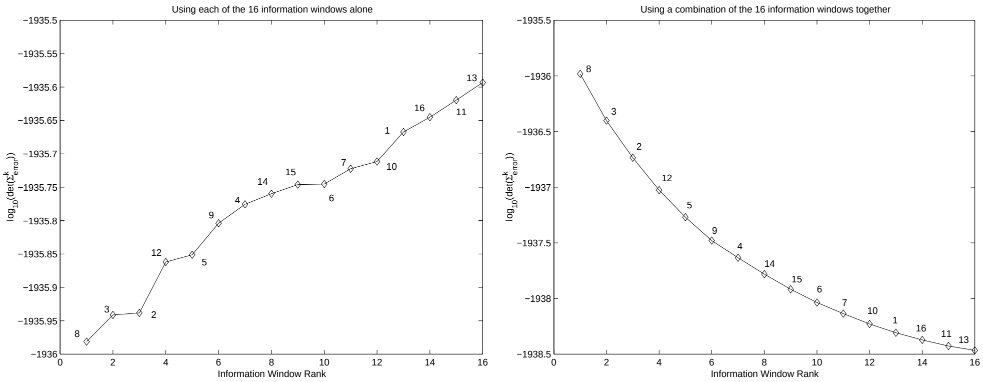

Fig. 4. Graphs of the information contained in each information window versus the window rank when using Simple Search (left) and Combinatorial Search (right). The numbers along the graph line are the window numbers.

In terms of computation time, Simple Search took an average of 6.2 seconds to rank the information windows while Combinatorial Search took an average of 63.9 seconds to determine the same information. The trade-off is accuracy versus computational power.

4.2 Window Results

We can divide our results into two main categories: Reconstruction using Information Windows and Position Estimation using Eigenwindows.

1. **Reconstruction using Information Windows**
   The reconstruction results obtained using information windows of size 8×8 pixels and omnidirectional images of 32×32 pixels in size are shown in Figure 5. Figure 5 (left) shows an omnidirectional image from

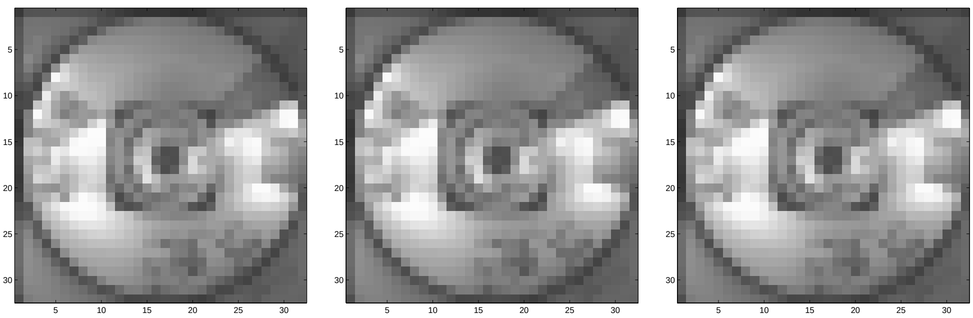

Fig. 5. Left: A 32×32 omnidirectional image acquired at runtime. Middle: Its reconstruction using the *most informative* information window. Right: Its reconstruction using all of the information windows. Each information window is 8×8 pixels in size.

   the a priori set, Figure 5 (middle) its reconstruction using only the *most informative* 8×8 information window and 5 (right) its reconstruction using all of the information windows. Reconstruction was achieved using equation (3). As can be seen from the images, a good reconstruction is obtained using only the best information window. This is an indication of the power of *Information Sampling*.
   Figure 6 shows the same experiment but using panoramic images as input data. Here the information windows were 8×10 pixels in size.
2. **Position Estimation using Eigenwindows**
   Having previously selected the best window using *Information Sampling*, we built a local appearance space using only the data selected with this window from each a priori image. This further compressed

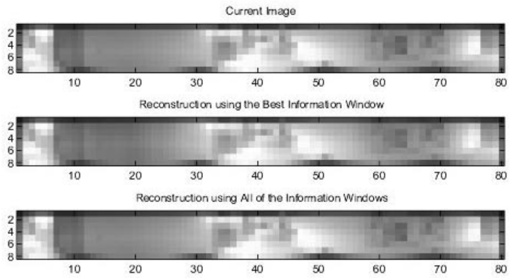

Fig. 6. Top: An 8×80 panoramic image acquired at runtime. Middle: Its reconstruction using the *most informative* information window. Bottom: Its reconstruction using all of the information windows. Each information window is 8×10 pixels in size.

our data to only approximately one thousandth of the original 128×128 image data. Successful position estimation has been achieved using windows as small as 4×4 pixels in size.
We projected *only* the selected windows from each image into the eigenspace. This is an improvement on previous approaches, where all windows first had to be projected. Thus, we were able to immediately reduce the ambiguity associated with projection. The images in Figure 7 show the results obtained using windows of 32×32 pixels in size.

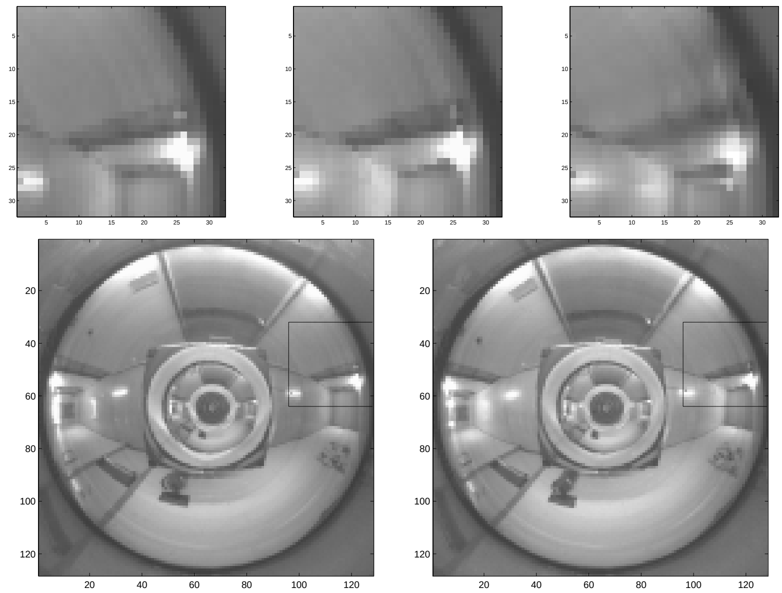

Fig. 7. Top: Close-up of the 32×32 information windows: unknown (left), closest (middle) and reconstructed (right). Bottom: The position of the unknown and closest images in their respective omnidirectional images.

The top row, from left to right shows the most relevant information window from an unknown image, its closest match from the a priori set of omnidirectional images and its reconstruction using equation (7). The bottom row, shows the information window in the unknown 128×128 image (left) and its closest match from the a priori set obtained by projecting only the most relevant information window (right). We note here that we could in principal, given enough computing power, use equation (3) to reconstruct a 128×128 *image* using only the most relevant *window*.

5 Conclusions and Future Work

In this paper we presented a statistical method termed *Information Sampling* which extracts, from a set of images, the data points containing the greatest amount of information. For computational reasons, *Infor-*

mation Sampling selected the points as *Information Windows*. We showed how this information could be used in reconstruction and presented two methods to rank the information windows from best to worst. We showed how to use only the best information to determine the qualitative position of a mobile robot within its environment. This position estimation technique was an improvement over existing approaches since we only needed to project the best data, thus reducing ambiguity.

Our future work will be directed towards extending the implementation of *Information Sampling*. In the short term the method shall be applied to the problem of image retrieval from a large database. In terms of mobile robotics, the possibility of rebuilding the local appearance space in real time, as a robot explores new regions of its environment shall be investigated.

Acknowledgements

This work was partly funded by the European Union RTD - Future and Emerging Technologies Project Number: IST-1999-29017, Omniviews.

References

1. M. J. Black and A. D. Jepson. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. *International Journal of Computer Vision*, 26(1):63–84, 1998.
2. V. Colin de Verdière and J. L. Crowley. Visual recognition using local appearance. In *5th European Conference on Computer Vision, (ECCV 1998)*, pages 640–654, Freiburg, Germany, June 1998.
3. R. Deriche and G. Giraudon. A computational approach for corner and vertex detection. *International Journal of Computer Vision*, 10(2):101–124, April 1993.
4. J. Gaspar, N. Winters, and J. Santos-Victor. Vision-based navigation and environmental representations with an omni-directional camera. *IEEE Transactions on Robotics and Automation*, 16(6):890–898, December 2000.
5. C. Harris and M. Stephens. A combined corner and edge detector. In *Alvey Vision Conference*, pages 147–151, 1988.
6. D. Jugessur and G. Dudek. Local appearance for robust recognition. In *Proceedings of the International Conference on Computer Vision and Pattern Recognition*, pages 834–839, Hilton Head Island, SC, USA, June 2000.
7. H. Murase, F. Kimura, F. Yoshimura, and Y. Miyake. An improvement of the autocorrelation matirx in pattern matching method and its application to hand printed 'hiragana'. *Transactions of IEICE*, J64-D(3):276–283, 1981.
8. H. Murase and S. K. Nayar. Visual learning and recognition of 3d objects from appearance. *International Journal of Computer Vision*, 14(1):5–24, January 1995.
9. K. Ohba and K. Ikeuchi. Detectibility, uniqueness and reliability of eigen windows for stable verification of partially occluded objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):1043–1048, September 1997.
10. D. Reisfeld, H. Wolfson, and Y. Yeshurun. Context free attention operators: the generalized symmetry transform. *International Journal of Computer Vision*, 14(2):119–130, March 1995.
11. M. Rendas and M. Perrone. Using field subspaces for on-line survey guidance. In *Proceedings of Oceans 2000*, Providence, RI, USA, June 2000.
12. A. P. Sage and J. L. Melsa. *Estimation Theory with Applications to Communications and Control*. McGraw-Hill, 1971.
13. C. Schmid and R. Mohr. Local grayvalue invariants for image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(5):530–535, May 1997.
14. L. Sirovich and M. Kirby. Low dimentional procedure for the characterization of human faces. *Journal of the Optical Society of America*, 4(3):519–524, 1987.
15. M. A. Turk and A. P. Pentland. Face recognition using eigenfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'91)*, pages 586–591, 1991.
16. N. Winters, J. Gaspar, and J. Santos-Victor. Omni-directional vision for robot navigation. In *Proceedings of the 1st International IEEE Workshop on Omni-directional Vision at CVPR 2000*, pages 21–28, Hilton Head Island, SC, USA, June 2000.
17. N. Winters and J. Santos-Victor. Omni-directional visual navigation. In *Proceedings of the 7th International Symposium on Intelligent Robotics Systems*, pages 109–118, Coimbra, Portugal, July 1999.

This article was processed using
the TEX macro package with SIRS2001 style