Robust Recovery of Eigenimages in the Presence of Outliers and Occlusions

Aleš Leonardis[1,2] and Horst Bischof[1]

[1] Department for Pattern Recognition and Image Processing, Technical University Vienna, Vienna, Austria
[2] Faculty of Computer and Information Science, University of Ljubljana, Ljubljana, Slovenia

The basic limitations of the current appearance-based matching methods using eigenimages are non-robust estimation of coefficients and inability to cope with problems related to occlusions and segmentation. In this paper we present a new approach which successfully solves these problems. The major novelty of our approach lies in the way how the coefficients of the eigenimages are determined. Instead of computing the coefficients by a projection of the data onto the eigenimages, we extract them by a hypothesize-and-test paradigm using subsets of image points. Competing hypotheses are then subject to a selection procedure based on the Minimum Description Length principle. The approach enables us not only to reject outliers and to deal with occlusions but also to simultaneously use multiple classes of eigenimages.

Keywords: appearance-based matching, principal component analysis, robust estimation, occlusion, discrete optimization.

1. Introduction and motivation

The appearance-based approaches to vision problems have recently received a renewed attention in the vision community due to their ability to deal with combined effects of shape, reflectance properties, pose in the scene, and the illumination conditions [Murase & Nayar, 1995b]. Besides, the appearance-based representations can be acquired through an automatic learning phase which is not the case with traditional shape representations. The approach has led to a variety of successful applications, e.g., illumination planning [Murase & Nayar, 1994], visual positioning and tracking of robot manipulators [Nayar et al., 1994], visual inspection [Yoshimura & Kanade, 1994], "image spotting" [Murase & Nayar, 1995a], and human face recognition [Turk & Pentland, 1991, Beymer & Poggio, 1995].

As stressed by its proponents, the major advantage of the approach is that both learning as well as recognition are performed using just two-dimensional brightness images without any low- or mid-level processing. However, there still remain various problems to be overcome since the technique rests on direct appearance-based matching [Murase & Nayar, 1995b]. The most severe limitation of the method in its present form is that it cannot handle the problems related to

- occlusion and

- segmentation.

The approach of modular eigenspaces [Pentland et al., 1994] tries to alleviate the problem of occlusion but does not solve it because the same limitation holds for each of the modular eigenspaces.

Moreover, the current approaches are also not robust, where the term *robustness* refers to the fact that the results remain stable in the presence of various types of noise and can tolerate a certain portion of outliers [Huber, 1981, Rousseuw & Leroy, 1987]. Robustness can be characterized by the concept of *breakdown point*, which is determined by the smallest portion of outliers in the data set at which the estimation procedure can produce an arbitrarily wrong estimate. For example, in current approaches even a single erroneous data point can cause an arbitrary wrong result, meaning that the breakdown point is 0%.

In this paper we present a new approach which successfully solves these problems. The major

novelty of our approach lies in the way how the coefficients of the eigenimages are determined. Instead of computing the coefficients by a projection of the data onto the eigenimages, we extract them by a hypothesize-and-test paradigm using *subsets* of image points. Competing hypotheses are then subject to a selection procedure based on the Minimum Description Length (MDL) principle. The approach enables us not only to reject outliers and to deal with occlusions but also to simultaneously use multiple classes of eigenimages.

The paper is organized as follows: We first review the basic concepts of the current appearance-based matching methods and point out the main limitations. In section 3 we present the basic steps of our method and outline the complete algorithm. The experimental results are shown in section 4: We first present some results on 1-D signals where the main steps of the algorithm can easily be visualized, and then we present the results on complex image data. We conclude with a summary and outline the work in progress.

2. Appearance-based matching

The appearance-based methods consist of two stages. In the first stage a set of images (templates), i.e., training samples, is obtained. These images usually encompass the appearance of a single object under different orientations [Yoshimura & Kanade, 1994], different illumination directions [Murase & Nayar, 1994], or multiple instances of a class of objects, e.g., faces [Turk & Pentland, 1991]. The sets of images are normally highly correlated. Thus, they can efficiently be compressed using Karhunen-Loève transform (i.e., PCA) [Anderson, 1958], resulting in a low-dimensional eigenspace.

In the second stage, given an input image, the recognition system projects parts of the input image, (i.e., subimages of the same size as training images), to the eigenspace. The recovered coefficients indicate the particular instance of an object and/or its position, illumination, etc. The process is usually sequentially applied to the entire image or some mechanisms are used to segment the subimages, e.g., motion.

We now introduce the notation. Let $\mathbf{y} = [y_1.$ $\ldots y_m]^T \in \mathbb{R}^m$ be an individual template, and

$\mathcal{Y} = \{\mathbf{y}_1 \ldots \mathbf{y}_n\}$ be a set of templates; throughout the paper a simple vector notation is used since the extension to 2-D is straightforward. To simplify the notation we assume $\mathcal{Y}$ to be normalized, having zero mean. Let $\mathbf{Q}$ be the covariance matrix of the vectors in $\mathcal{Y}$; we denote the eigenvectors of $\mathbf{Q}$ by $\mathbf{e}_i$, and the corresponding eigenvalues by λ_i. We assume that the number of templates n is much smaller than the number of elements m in each template, thus an efficient algorithm based on SVD can be used to calculate the first n eigenvectors [Murase & Nayar, 1995b]. Since the eigenvectors form an orthogonal basis system, $< \mathbf{e}_i, \mathbf{e}_j > = 1$ when $i = j$ and 0 otherwise, where $<>$ stands for a scalar product. We assume that the eigenvectors are ordered in the descending order with respect to the corresponding eigenvalues λ_i. Then, depending on the correlation among the templates in $\mathcal{Y}$, only p, $p < n$, eigenvectors are needed to represent the $\mathbf{y}_i$ to a sufficient degree of accuracy as a linear combination of eigenvectors $\mathbf{e}_i$

$$\tilde{\mathbf{y}} = \sum_{i=1}^{p} a_i(\mathbf{y})\mathbf{e}_i \ . \tag{1}$$

We call the space spanned by the first p eigenvectors the *eigenspace*.

To recover the parameters a_i during the matching stage, a data vector $\mathbf{x}$ is projected onto the eigenspace

$$a_i(\mathbf{x}) = < \mathbf{x}, \mathbf{e}_i > = \sum_{j=1}^{m} x_j e_{i_j} \quad 1 \leq i \leq p \ . \tag{2}$$

$\mathbf{a}(\mathbf{x}) = [a_1(\mathbf{x}) \ldots a_p(\mathbf{x})]^T$ is the point in the eigenspace obtained by projecting $\mathbf{x}$ onto the eigenspace. Let us call the $a_i(\mathbf{x})$ coefficients of $\mathbf{x}$. The reconstructed data vector $\tilde{\mathbf{x}}$ can be written as

$$\tilde{\mathbf{x}} = \sum_{i=1}^{p} a_i(\mathbf{x})\mathbf{e}_i \ . \tag{3}$$

It is well known that PCA is among all linear transformations the one which is optimal with respect to the reconstruction error $||\mathbf{x} - \tilde{\mathbf{x}}||^2$.

(a) Two eigenfunctions   (b) Test signal

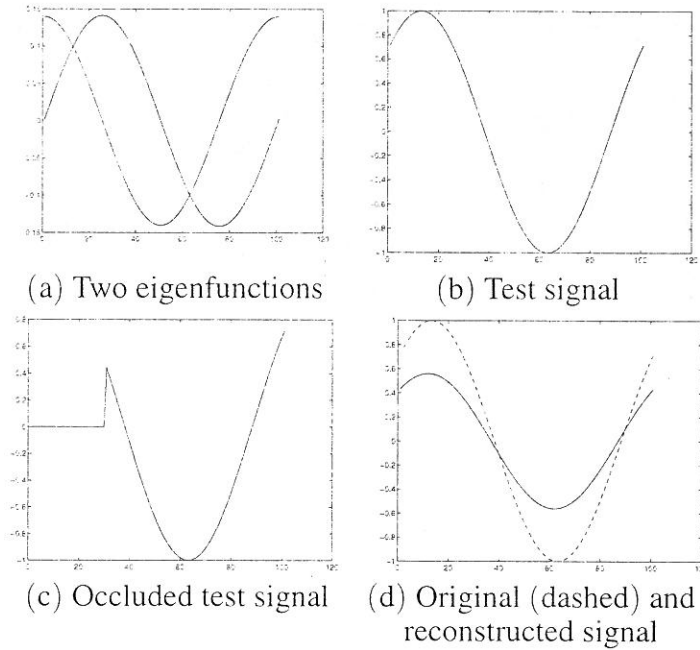(c) Occluded test signal   (d) Original (dashed) and reconstructed signal

Fig. 1. Demonstration of the occlusion using the standard approach for calculating the coefficients a_i

2.1. Weaknesses of current appearance-based matching

In this section we analyze some of the basic limitations of the current appearance-based matching methods and illustrate them with a few examples. Namely, the way how the coefficients a_i are calculated poses a serious problem in the case of outliers and occlusions.

Outliers Let us suppose that x_j in Eq. (2) is corrupted by δ. Then, $\hat{a}_i = a_i + \delta e_{i_j}$. It follows that $\|\hat{\mathbf{a}} - \mathbf{a}\|$ can get arbitrarily large, just by changing a single component x_j. This proves that the method is non-robust with the breakdown point 0%.

Occlusion Similarly one can analyze the effect of occlusion. Suppose that $\hat{\mathbf{x}} = [x_1, \ldots, x_r, 0, \ldots, 0]^T$ is obtained by setting last $m - r$ components of $\mathbf{x}$ to zero; a similar analysis holds when some of the components of $\mathbf{x}$ are set to some other values, which, for example, happens in the case of occlusion by another object. Then

$$\hat{a}_i = \hat{\mathbf{x}}^T \mathbf{e}_i = \sum_{j=1}^{r} x_j e_{i_j} . \qquad (4)$$

The error we make in calculating a_i is

$$(a_i(\mathbf{x}) - \hat{a}_i(\hat{\mathbf{x}})) = \sum_{j=r+1}^{m} x_j e_{i_j} . \qquad (5)$$

Similarly, the additional error caused by occlusion is

$$\| \sum_{i=1}^{p} (\sum_{j=r+1}^{m} x_j e_{i_j}) \mathbf{e}_i \|^2 . \qquad (6)$$

This error is not localized at the occluded part but spreads over the whole vector $\tilde{\mathbf{x}}$.

Let us demonstrate the effect of occlusion on a simple 1-D example.

Fig. 1a shows the two eigenfunctions, obtained as a result of training on a set of trigonometric functions (Fig. 4, section 4). Fig. 1b shows a test-signal taken from the set of training signals. The signal can be exactly described by the coefficients $\mathbf{a} = [5.0042, 5.0470]^T$. Fig. 1c shows an occluded signal, where the first 30 elements have been set to 0. Using Eqs. (2) and (3) we get the signal shown in Fig. 1d, with the calculated coefficient vector $\mathbf{a} = [2.5658, 3.0509]^T$. This simple example nicely demonstrates the consequences of the occlusion. Of course, the same applies to the 2-D case as shown in Fig. 2.

The problems that we have discussed arise because the complete set of data $\mathbf{x}$ is required to calculate a_i in Eq. (2). Therefore the method is
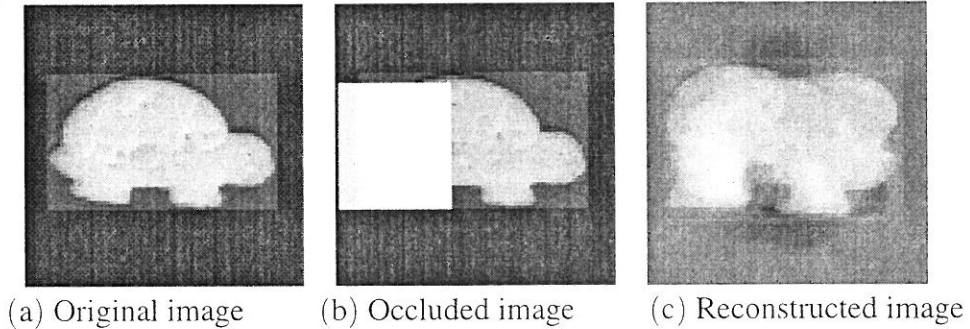
(a) Original image   (b) Occluded image   (c) Reconstructed image

Fig. 2. Demonstration of the occlusion using the standard approach for calculating the coefficients a_i

sensitive to partial occlusions, to data containing noise and outliers, and to changing backgrounds. In the next section we explain our new approach which has been designed to overcome precisely this type of problems.

3. Our approach

The major novelty of our approach lies in the way how the coefficients of the eigenimages are determined. Instead of computing the coefficients by a projection of the data onto the eigenimages, we extract them by a robust hypothesize-and-test paradigm using only *subsets* of image points. Competing hypotheses are then subject to a selection procedure based on the Minimum Description Length principle. More specifically, our approach, which we present in the following subsections, consists of four main steps: hypotheses generation, selection, fitting, and a final selection.

3.1. Generating hypotheses

Let us first start with a simple observation. In order to calculate the coefficients a_i (Eq. 2) we need only p points $\mathbf{r} = (r_1, \ldots, r_p)$. This is sufficient to compute the coefficients a_i by simply solving the following system of linear equations:

$$x_{r_i} = \sum_{j=1}^{p} a_j(\mathbf{x})e_{j r_i} \quad 1 \leq i \leq p \ . \quad (7)$$

The coefficients a_i are then used to create a hypothesis $\tilde{\mathbf{x}}$, which is tested by the backprojection using Eq. (3). The test gives us an error

vector $\vec{\xi} = (\mathbf{x} - \tilde{\mathbf{x}})$. The points which are within an error margin Θ form a set of *compatible* points according to hypothesis $\tilde{\mathbf{x}}$ generated from $\mathbf{a}$. We treat a hypothesis as acceptable if it contains a minimal number of compatible points. This condition can really be kept minimal since the selection procedure will reject the false positives. The accepted hypothesis is characterized by the coefficient vector $\mathbf{a}$, the error vector $\vec{\xi}$, and the domain of the compatible points $D = \{j | \xi_j^2 < \Theta\}, s = |D|$.

The considerations regarding Eq. (7) hold if we take into account all eigenvectors, i.e., $p = n$, and if there is no noise in the data x_{r_i}. However, if we approximate each template only by a linear combination of a subset of eigenfunctions and there is also noise present in the data, then Eq. (7) can no longer be used, but rather we have to solve an overconstrained system of equations in the least squares sense using k data points ($p < k << m$). Thus we seek the solution vector $\mathbf{a}$ which minimizes

$$E(\mathbf{r}) = \sum_{i=1}^{k} (x_{r_i} - \sum_{j=1}^{p} a_j(\mathbf{x})e_{j r_i})^2 \ . \quad (8)$$

Of course, the minimization of Eq. (8) can only produce correct values for coefficient vector $\mathbf{a}$, if the set of points r_i does not contain outliers, i.e, not only extreme noisy points but also points belonging to different backgrounds or some other templates due to occlusion. Therefore, the solution has to be sought in a robust manner. In particular, we randomly select a set of points, e.g., $k = 5p$ and then iteratively, based on the error distribution, reduce their number to approximately $3p$ points, which give us the final

solution of the minimum least squares problem, Eq. (8). To efficiently search for subsets of data points to initialize the hypotheses one can use a data-driven masking technique [Stricker & Leonardis, 1995].

Once we have the coefficient vector **a**, we can evaluate the hypothesis, both from the point of view of the error, where we can expect for good matches an error of $\sum_{i=p+1}^{n} \lambda_i$ on the average, and from the number of compatible points. However, one can not expect that every initial randomly chosen set of points will produce a good hypothesis if there is one, despite the robust procedure. Thus, to further increase the robustness of the hypotheses generation step, i.e., increase the probability of detecting a correct hypothesis if there is one, we initiate a number of trials. This leads to a possibly redundant set of accepted hypotheses, which is then resolved by the selection procedure.

3.2. Selection

The set of hypotheses which has been generated is usually highly redundant. Thus, the selection procedure has to select a subset of "good" hypotheses and reject the superfluous ones. To achieve this, we utilize the Minimum Description Length principle, which leads to the minimization of an objective function encompassing the information about the competing hypotheses [Leonardis *et al.*, 1995]. The objective function has the following form:

$$F(\mathbf{h}) = \mathbf{h}^T \mathbf{C} \mathbf{h} = \mathbf{h}^T \begin{bmatrix} c_{11} & \cdots & c_{1R} \\ \vdots & & \vdots \\ c_{R1} & \cdots & c_{RR} \end{bmatrix} \mathbf{h} \ . \tag{9}$$

Vector $\mathbf{h}^T = [h_1. h_2. \ldots . h_R]$ denotes a set of hypotheses, where h_i is a *presence-variable* having the value 1 for the presence and 0 for the absence of the hypothesis i in the resulting description. The diagonal terms of the matrix **C** express the cost-benefit value for a particular hypothesis i

$$c_{ii} = \mathrm{K}_1 s_i - \mathrm{K}_2 ||\vec{\xi_i}||_{D_i} - \mathrm{K}_3 N_i \ , \tag{10}$$

where s_i is the number of compatible points, $||\vec{\xi_i}||_{D_i}$ is the error over the domain D_i, and N_i is

the number of coefficients (eigenvectors). The coefficients K_1, K_2, and K_3, which can be determined automatically [Leonardis *et al.*, 1995], adjust the contribution of the three terms. K_1 is related to the average cost of describing a data point (in bits), K_2 is related to the average cost of specifying the error (follows from Θ), and K_3 is related to the average cost of specifying a coefficient of an eigenvector. Due to the nature of the problem, i.e., finding the maximum of the objective function, only the relative ratios between the coefficients play a role, e.g., $\mathrm{K}_2/\mathrm{K}_1$ and $\mathrm{K}_3/\mathrm{K}_1$.

The off-diagonal terms handle the interaction between the overlapping hypotheses

$$c_{ij} = \frac{-\mathrm{K}_1 |D_i \cap D_j| + \mathrm{K}_2 \xi_{ij}}{2}$$
$$\xi_{ij}^2 = \max(\sum_{D_i \cap D_j} \vec{\xi_i}. \sum_{D_i \cap D_j} \vec{\xi_j}) \ . \tag{11}$$

where D_i denotes the domain of the i-th hypothesis and $\sum_{D_i \cap D_j} \vec{\xi_i}$ denotes the sum of squared errors of the i-th hypothesis over the intersection of the two domains D_i, D_j.

The objective function takes into account the interaction between different hypotheses which may be completely or partially overlapped. However, we consider only the pairwise overlaps in the final solution. From the computational point of view, it is important to notice that the matrix **C** is symmetric, and depending on the overlap, it can be sparse or banded. All these properties of the matrix **C** can be used to reduce the computations needed to calculate the value of $F(\mathbf{h})$.

We have now formulated the problem of selection in such a way that its solution corresponds to the global extremum of the objective function. Maximization of the objective function $F(\mathbf{h})$ belongs to the class of combinatorial optimization problems (quadratic Boolean problem). Since the number of possible solutions increases exponentially with the size of the problem, it is usually not tractable to explore them exhaustively. Thus the exact solution has to be sacrificed to obtain a practical one. Various methods have been proposed for finding a "global extreme" of a class of nonlinear objective functions. Among these methods

are winner-takes-all strategy, simulated annealing, microcanonical annealing, mean field annealing, Hopfield networks, continuation methods, and genetic algorithms [Cichocki & Unbehauen, 1993]. We are currently using two different methods for optimization. One is a simple greedy algorithm, which is computationally very efficient, and the other one is Tabu search [Glover & Laguna, 1993, Stricker & Leonardis, 1995]. Tabu search is computationally a little more demanding but it provides consistently better results than the greedy algorithm.

3.3. Fitting

The selected (also the generated) hypotheses are based only on $3p$ points, from which the coefficients have been computed. In order to increase the accuracy of the coefficient vector **a**, we use a modified (robustified) least squares approach [Stricker, 1994]. The principle is simple: Based on the statistics of the errors of *all* compatible points, a decision is made (again related to the error margin Θ) which data points should participate in a standard least squares computation of **a**. This is repeated until the convergence is reached. Since for the selected
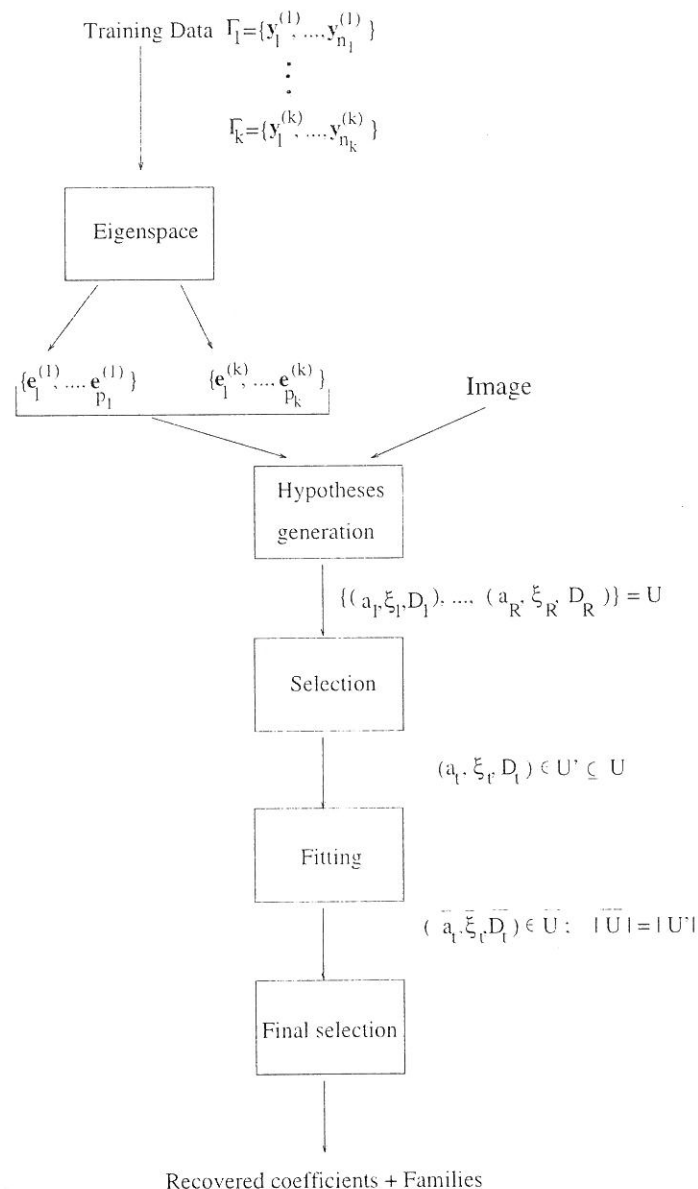

Training Data $\Gamma_1 = \{y_1^{(1)}, \ldots, y_{n_1}^{(1)}\}$

$$\vdots$$

$$\Gamma_k = \{y_1^{(k)}, \ldots, y_{n_k}^{(k)}\}$$

Eigenspace

$$\{e_1^{(1)}, \ldots, e_{p_1}^{(1)}\} \qquad \{e_1^{(k)}, \ldots, e_{p_k}^{(k)}\} \qquad \text{Image}$$

Hypotheses generation

$$\{(a_1, \xi_1, D_1), \ldots, (a_R, \xi_R, D_R)\} = U$$

Selection

$$(a_t, \xi_t, D_t) \in U' \subseteq U$$

Fitting

$$(\bar{a}_t, \bar{\xi}_t, \overline{D_t}) \in \bar{U}: \quad |\bar{U}| = |U'|$$

Final selection

Recovered coefficients + Families

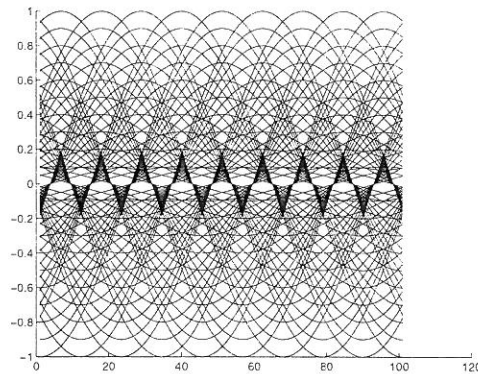*Fig. 3.* A schematic diagram outlining the complete algorithm.

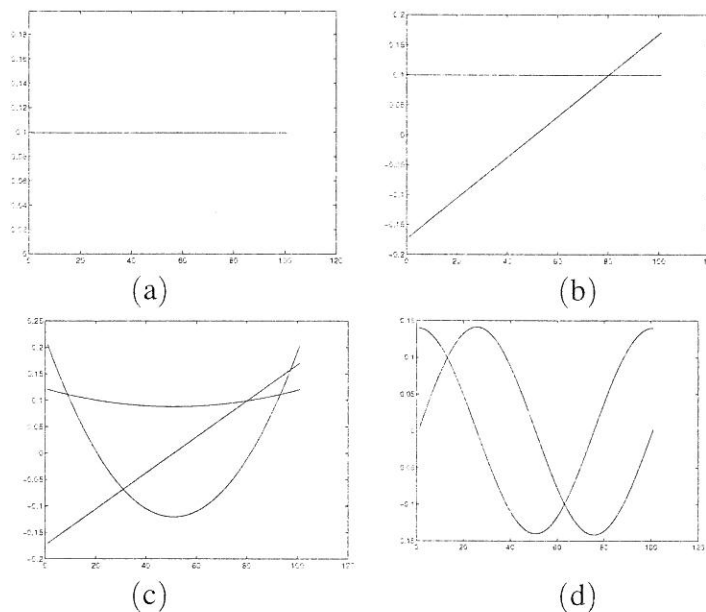*Fig. 4.* Training samples: Trigonometric function.

(a)

(b)

(c)

(d)

Fig. 5. 1-D eigensignals: (a) constant, (b) linear, (c) quadratic, and (d) trigonometric.

hypotheses the initial estimates of **a** are usually quite accurate, only a few iterations are necessary for the process to converge.

3.4. Final selection

Since the previous fitting step might have altered some of the hypotheses, a final selection is in some cases needed. The selection is performed in the same manner as explained in section 3.2. Note that the last two steps are computationally not very expensive and cost only a small fraction of time compared to the first

two steps, since they are applied only to a small subset of hypotheses.

3.5. Complete algorithm

The complete algorithm is shown in Fig. 3. One should note that we can simultaneously deal with multiple eigenspaces $\Gamma_1, \ldots, \Gamma_k$ of, for example, different objects (we call them *families*). Everything remains as described, except that at each location all of the families are used to initiate the hypotheses. The selection procedure then reasons among different hypotheses, possibly belonging to different families, and selects those that better explain the data.
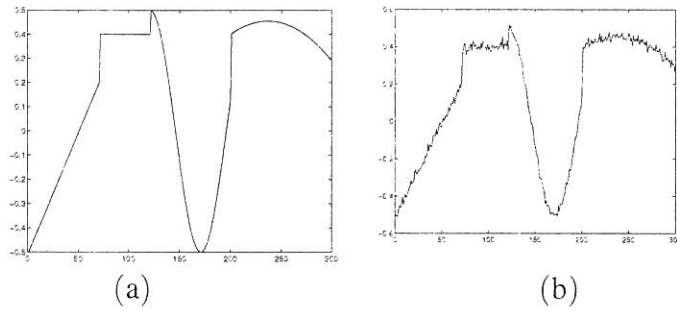
(a)   (b)

Fig. 6. 1-D test-signals: (a) noiseless, (b) signal with added Gaussian noise.

(a) hypotheses generation   (b) selection

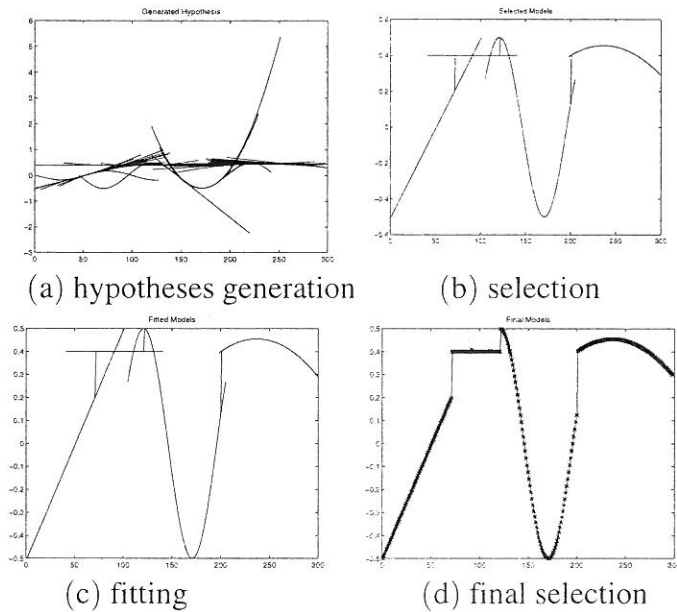(c) fitting   (d) final selection

Fig. 7. The four steps of our algorithm shown on a 1-D signal. In Figs. (a–c) the complete extent of hypotheses is depicted while in Fig. (d) only the compatible points are plotted.

4. Experimental results

4.1. 1-D case

In order to test the algorithm we have first applied it to 1-D functions. We started with generating four families of functions:

- $y = a$

- $y = bx + c$

- $y = dx^2 + ex + f$

- $y = g \cos(x + h)$

For each of these families 200 training examples have been generated by systematically varying their parameters $(a \dots h)$. Fig. 4 shows some of the training samples for the trigonometric family.

These samples are then used to calculate the four eigenspaces (we have used all eigenvectors with $\lambda_i \neq 0$). Fig. 5 shows the obtained eigenvectors.

We have tested our algorithm on various test-signals like the ones shown in Fig. 6. Note that these test-signals contain only portions of the original functions and are therefore not recoverable by the standard eigenspace approach (as has been demonstrated in Fig. 1).

Fig. 7 shows the four steps of our algorithm applied to the test-signal depicted in Fig 6a. The members of all four families have been perfectly recovered, as depicted in Fig. 7d. One can also see from this example that the fitting step did
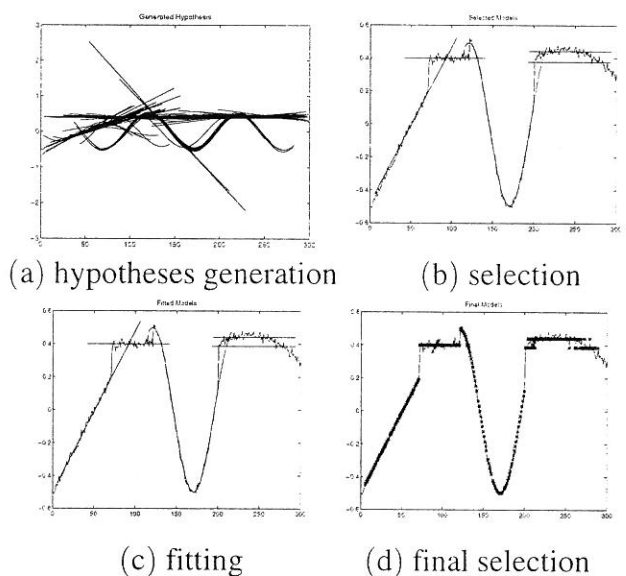
(a) hypotheses generation   (b) selection

(c) fitting   (d) final selection

Fig. 8. The four steps of our algorithm shown on a noisy 1-D signal. In Figs. (a–c) the complete extent of hypotheses is depicted while in Fig. (d) only the compatible points are plotted.
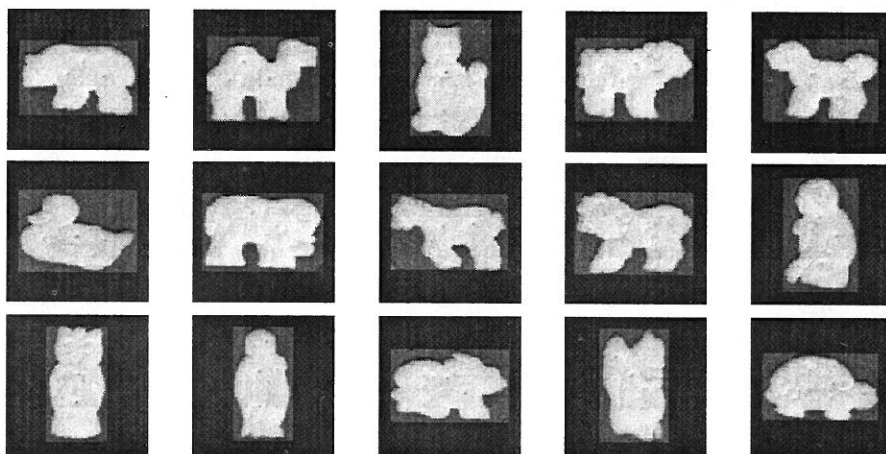

Fig. 9. Set of training templates (a family of biscuit-animals).

not change the hypotheses very much, therefore the final selection has not changed anything.

Fig. 8 shows the four steps of the algorithm in the case of a noisy signal. In addition we have not included the family of quadratic functions. Therefore one gets for the quadratic part an approximation consisting of two constant functions. This example demonstrates that our algorithm not only successfully deals with partial signals but also tolerates noise.

Further extensive experimentation on a variety of different signals has confirmed the robustness of the approach.

4.2. 2-D case

To demonstrate the proposed method on real images we took three sets of images. The first set contains a family of "biscuit-animals" (see Fig. 9) while the other two sets encompass the appearance of the "duck" and the "camel" under thirty-two orientations. The images in all three sets are correlated to a large degree, resulting in a small number of eigenvectors (eigenimages). Figs. 10, 11, and 12 show the first 10 eigenimages of each family. For the experiments we have used 8 eigenvectors for the biscuit-animals and 16 for the camel and duck family, respectively.
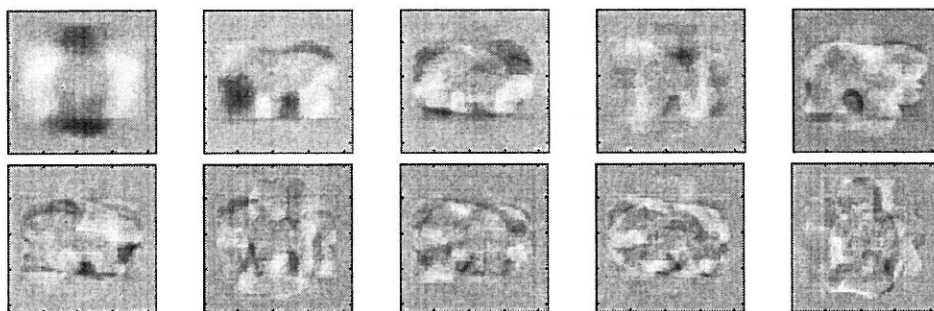
Fig. 10. First ten eigenvectors corresponding to the family of animals.
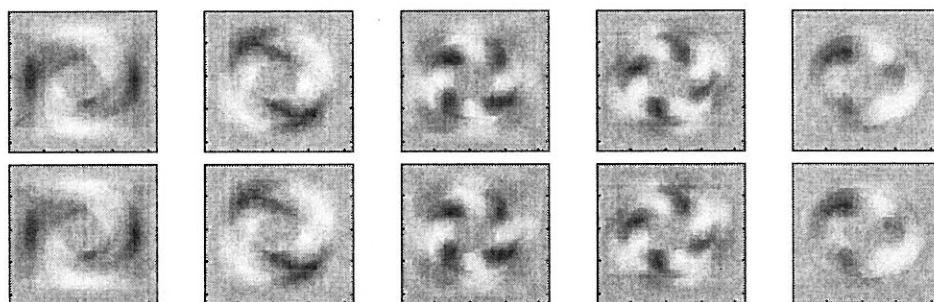

Fig. 11. First ten eigenvectors corresponding to the family of rotated camels.
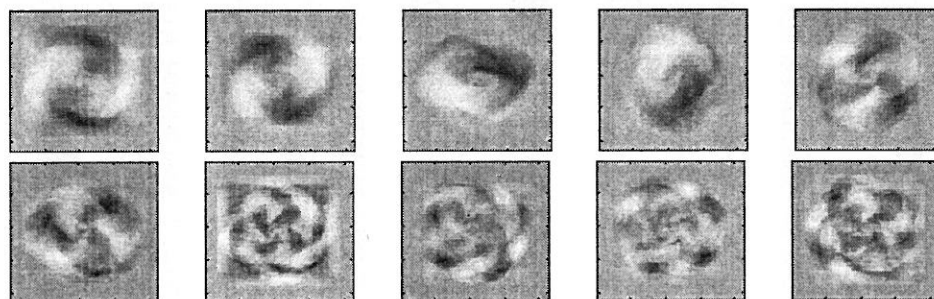

Fig. 12. First ten eigenvectors corresponding to the family of rotated ducks.

Figures 13 and 14 show two of the test images we have used. In both images, there is a highly cluttered background. In addition, in Fig. 14 we have partially occluded objects.

Cluttered background   Fig. 15a shows a portion of the scene from Fig. 13 containing a camel. The algorithm outputs the family (i.e., the camel-family) which best describes the data, and the coefficients from which the orientation is identified. For illustration purpose, Fig. 15b shows the backprojection of the coefficients that have been recovered by our algorithm.

Occluded objects   Fig. 16b shows the occluded turtle from Fig. 16a, and Fig. 16c gives the result obtained by our algorithm. The result is

perfect in terms of the recovered coefficients, demonstrating that the method can deal with a considerable amount of occlusions.

Figure 17a shows a portion of the scene from Fig. 14, depicting a part of a camel (containing approx. 30% of the pixels of the original camel) on a cluttered background. The recovered information suffices to determine the family of the object (camel), and its orientation. Again, to illustrate, Fig. 17b shows the backprojection of the coefficients that have been recovered by the algorithm. This example demonstrates that our algorithm is very robust in dealing both with occlusions and cluttered background. The algorithm also correctly detected other animals in the image while in the areas with no animals, no false positives have been reported after the final selection.

Fig. 13. A scene containing some of the templates on a complex background.

Fig. 14. A scene containing some of the templates (some of them being partially occluded) on a complex background.

5. Conclusion

In this paper we have presented a novel approach which enables the appearance-based matching techniques to successfully cope with outliers, cluttered background, and occlusions.

The method exploits several techniques, e.g., robust estimation and hypothesize-and-test paradigm, which combined together in a general framework achieve the goal.

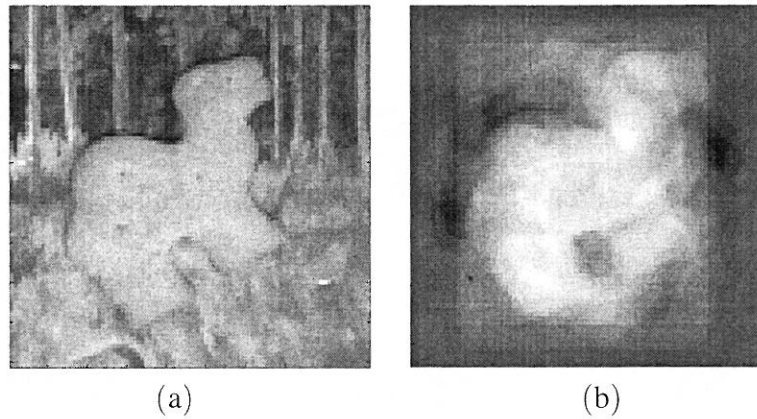The applications of the proposed approach are numerous. Basically everything that can be

Fig. 15. (a) A detail from the original image (Fig. 13) and (b) the reconstructed camel.
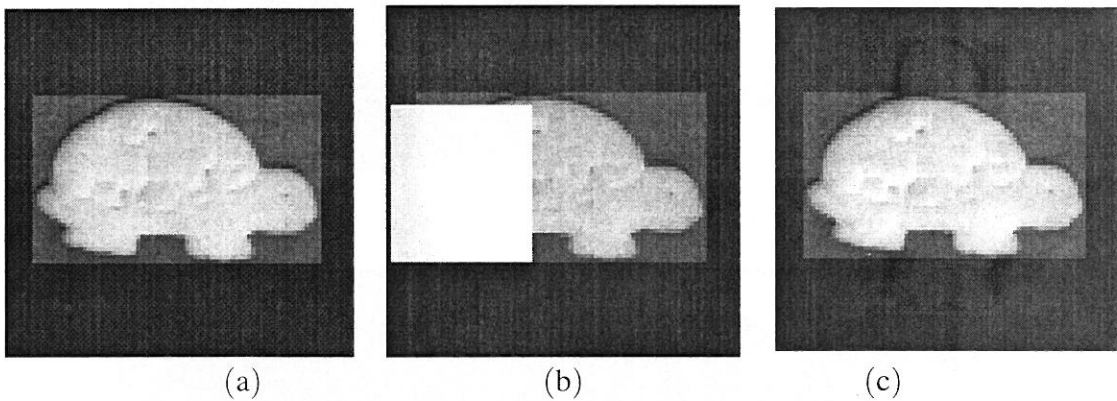

Fig. 16. (a) Original image, (b) occluded turtle, and (c) the reconstructed turtle, clearly being unaffected by the occlusion (compare with figure 2).
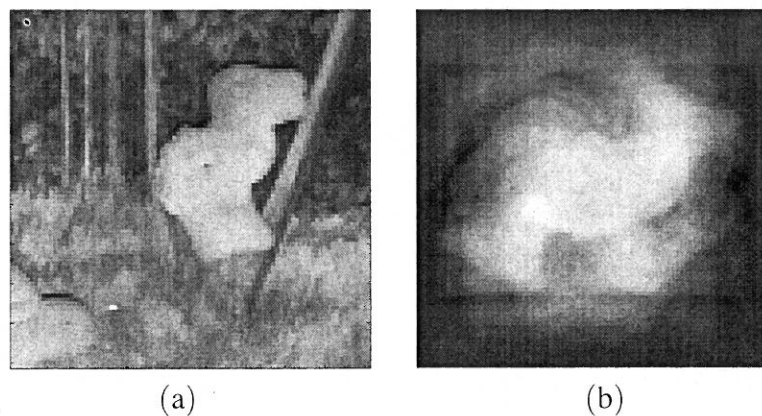

Fig. 17. (a) A detail from the original image (Fig. 14) and (b) the reconstructed camel.

performed with the classical appearance-based methods can also be achieved within the framework of our approach, only more robustly and in more complex scenes. Besides, when the proposed method is used on the type of images that can also be dealt with by the classical methods (free of outliers, occlusion, with a predefined background), one can expect significant computational gains since in our case, the coefficient vector is calculated directly rather than through the projection of the image data onto the eigenspace.

It is interesting to note that the basic steps of the proposed algorithm, namely hypotheses generation, selection, fitting, and the final selec-

tion, are the same as in ExSel+.+ [Stricker & Leonardis, 1995], which deals with robust extraction of *analytical* parametric functions from various types of data. Therefore, the method described in this paper can also be seen as an extension of ExSel++ to *learnable classes* of parametric models.

Our current work is directed towards further increasing the robustness of the recognition by exploiting the constraints that come from the fact that we are effectively working in a discrete coefficient space rather than in a continuous one. Namely, so far, we have been calculating the coefficients as they were continuous variables. We are also exploiting the possibility of incorporating the approach in a multiresolution framework.

Acknowledgements

References

[Anderson, 1958] T. W. ANDERSON, *An Introduction to Multivariate Statistical Analysis*. Wiley, New York, 1958.

[Beymer & Poggio, 1995] D. BEYMER & T. POGGIO, Face recognition from one example view. In *Proceedings of the 5th ICCV'95*, (1995), pp. 500–507. IEEE Computer Society Press.

[Cichocki & Unbehauen, 1993] A. CICHOCKI & R. UNBEHAUEN, *Neural Networks for Optimization and Signal Processing*. Wiley, New York, 1993.

[Glover & Laguna, 1993] F. GLOVER & M. LAGUNA, Tabu search. In *Modern Heuristic Techniques for Combinatorial Problems*. C. R. REEVES, (ed), (1993), pp. 70–150. Blackwell Scientific Publications.

[Huber, 1981] P. J. HUBER, *Robust Statistics*. Wiley, New York, 1981.

[Leonardis et al., 1995] A. LEONARDIS, A. GUPTA, & R. BAJCSY, Segmentation of range images as the search for geometric parametric models. *International journal of computer vision*, **14**(3), (1995), 253–277.

[Murase & Nayar, 1995a] H. MURASE & S. K. NAYAR, Image spotting of 3D objects using parametric eigenspace representation. In *Proceedings of the 9th Scandinavian Conference on Image Analysis*, (1995), G. BORGEFORS, (ed), Vol. 1, pp. 323–332.

[Murase & Nayar, 1995b] H. MURASE & S. K. NAYAR, Visual learning and recognition of 3-D objects from appearance. *International Journal of Computer Vision*, **14**, (1995), 5–24.

[Murase & Nayar, 1994] H. MURASE & S. K. NAYAR, Illumination planning for object recognition using parametric eigenspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **16**(12), (1994), 1219–1227.

[Nayar et al., 1994] S. K. NAYAR, H. MURASE, & S. A. NENE, Learning, positioning, and tracking visual appearance. In *Proceedings of IEEE International Conference on Robotics and Automation*, (1994).

[Pentland et al., 1994] A. PENTLAND, B. MOGHADDAM, & T. STRANER, View-based and modular eigenspaces for face recognition. Technical report 245, MIT Media Laboratory, 1994.

[Rousseuw & Leroy, 1987] P. J. ROUSSEUW & A. M. LEROY, *Robust Regression and Outlier Detection*. Wiley, New York, 1987.

[Stricker, 1994] M. STRICKER, A new approach for robust ellipse fitting. In *Proceedings of the Third International Conference on Automation, Robotics and Computer Vision*, (1994), Vol. 2, pp. 940–945.

[Stricker & Leonardis, 1995] M. STRICKER & A. LEONARDIS, ExSel++: A general framework to extract parametric models. In *Proceedings of the 6th International Conference on Computer Analysis of Images and Patterns, CAIP'95*, (1995), V. HLAVÁČ & R. ŠÁRA, (eds), pp. 90–97. Lecture Notes in Computer Science, Vol. 970. Springer, Prague, Czech Republic.

[Turk & Pentland, 1991] M. TURK & A. PENTLAND, Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, **3**(1), (1991), 71–86.

[Yoshimura & Kanade, 1994] S. YOSHIMURA & T. KANADE, Fast template matching based on the normalized correlation by using multiresolution eigenimages. In *Proceedings of IROS'94*, (1994), pp. 2086–2093.

Contact address:

Horst Bischof
Department for Pattern Recognition and Image Processing
Technical University Vienna
Treitlstraße 3/1832, A-1040 Vienna, Austria
phone: +43 (1) 58801 4478
fax: +43 (1) 5054 668
e-mail: {ales,bis}@prip.tuwien.ac.at

Aleš Leonardis
Computer Vision Laboratory
Faculty of Computer and Information Science
University of Ljubljana
Tržaška 25, 1000 Ljubljana, Slovenia
phone: +386 (61) 1768 381
fax: +386 (61) 1264 630
e-mail: ales.leonardis@fri.uni-lj.si

ALEŠ LEONARDIS received his B.S. degree and M.S. degree in electrical engineering from the University of Ljubljana, Slovenia in 1985 and 1988, respectively. From 1988–1991 he was a visiting researcher in the General Robotics and Active Sensory Perception Laboratory at the University of Pennsylvania. He received his Ph.D. in computer science from the University of Ljubljana in 1993, where he was elected Docent in 1994. His research interests include robust methods for computer vision, 3–D scene interpretation, recognition, and learning.

HORST BISCHOF received his M.S. and Ph.D. degrees in computer science from the Tehnical University Vienna, Austria in 1990 and 1993, respectively. From 1989–1991 he was at Remote Sensing Department at University of Agriculture. Since 1991 he is with the Department for Pattern Recognition and Image Processing at the Tehnical University Vienna. His research interests include theory and applications of neural networks for computer vision, learning, and recognition.