

Low-dimensional representation of faces in higher dimensions of the face space

Alice O'Toole*, Hervé Abdi*†, Kenneth, A. Deffenbacher** and Dominique Valentin*

*School of Human Development, The University of Texas at Dallas, Richardson, TX 75083-0688, †Université de Bourgogne, Dijon 21004 France; **The University of Nebraska at Omaha.

Received February 25 1992; accepted August 14, 1992; revised manuscript received September 15, 1992

Faces can be represented efficiently as a weighted linear combination of the eigenvectors of a covariance matrix of face images. It has also been shown [J. Opt. Soc. Am. 4, 519-524 (1987)] that identifiable faces can be made by using only a subset of the eigenvectors, i.e., those with the largest eigenvalues. This low-dimensional representation is optimal in that it minimizes the squared error between the representation of the face image and the original face image. The present study demonstrates that, whereas this low-dimensional representation is optimal for identifying the physical categories of face, like sex, it is not optimal for recognizing the faces (i.e., discriminating known from unknown faces). Various low-dimensional representations of the faces in the higher dimensions of the face space (i.e., the eigenvectors with smaller eigenvalues) provide better information for face recognition.

1. INTRODUCTION

Computational models of face recognition must address several difficult problems. First, faces must be represented in a way that makes available information that is sufficient to distinguish a particular face from all other faces. Faces pose a particularly difficult problem in this respect because all faces are similar to one another in that they contain the same set of features (e.g., eyes, nose, mouth) arranged in roughly the same configuration. Early attempts to model face recognition¹ generally attempted to use a geometric coding in which measurements of the relations between facial features were coded and used for recognition purposes. These features, however, discard important information about the texture and shape of the face² and about subtle variations in the form and configuration of the facial features³. Recent research on computational modeling of faces has found it useful to employ a simpler representation of faces that consists of a normalized pixel-based representation of

faces. This representation has been used to code complete faces³⁻⁹ partial faces¹⁰, and the subcomponents of faces². Whereas these approaches do not create an invariant three-dimensional representation, when they are used with full faces, they provide detailed information that is useful for distinguishing faces both at the configural level (i.e., the relationship between features) and at the featural level. Thus, geometric representations are coded implicitly, but texture and detailed shape information are also preserved. These pixel-based codings are generally used in conjunction with models that represent faces, explicitly or implicitly, as a weighted combination of eigenvectors.^{3,4,9,10} We discuss these models below in the context of the tasks to which they have been applied.

The second major problem that computational models must address is that effective computational modeling of faces requires not only the ability to recognize faces but also the ability to process faces in a variety of other ways. Because the terms to be used

for these other tasks have several meanings in different disciplines and in everyday language, we start by offering a definition of each term as it is used in this paper. "Recognition" is the categorization of a face as known or unknown. Whereas this task can be viewed as a categorization, it is important to note that this categorization is based not on any physical similarities between faces within a category but rather on a particular instance of experience with the face. In other words, a particular familiar face may be more physically similar to a particular unfamiliar face than to any other familiar face. A final point to note is that the definition of recognition that we use in this paper does not assume that the face is identifiable to the model as an instance of a particular person but rather assumes only that familiar or learned faces can be distinguished from unfamiliar or unlearned faces by the model. This distinction also holds in human memory. We are often absolutely certain that we recognize a person (or know the person from somewhere) but are completely unable to recall any other information about the person, such as their name or the place with which we associate them. Generally this happens when the person is seen out of context. For example, when we run into a local grocery store clerk at a movie theater, we know the face is familiar but are unable to access who the person is or where we have seen him/her before.

Beyond recognition, it is also useful to be able to extract physically based categorical information from unknown or unfamiliar faces. We are able to look at any unfamiliar face and make judgments about the sex, race, and approximate age of the person pictured. For lack of a better term, we will call this task categorization of visually derived semantic information after Bruce and Young's psychological model of human face processing.¹¹ In contrast to recognition, categorization in terms of visually derived semantic information is presumably based on the fact that faces within a particular category are more physically similar to one another, in some way, than are faces between categories. Therefore, e.g., female faces should be more similar to one another in some respect than female faces are to male faces.

The final distinction we make is between the kind of information that is needed for recognition and the kind of information that provides a best likeness of a face. Human face recognition works reasonably

well even in suboptimal conditions. We can recognize faces from blurry photographs and from photographs revealing only partial information about the face. Nonetheless we can also judge how good a likeness a particular picture is of a face and can determine the degree to which additional information improves or decreases the likeness.

The present studies use a principal components approach toward quantifying and describing the information that is most useful for these tasks. In this approach faces are represented as a weighted combination of eigenvectors that are extracted from a covariance matrix derived from a set of faces. The eigenvectors act as features in some traditional senses of the word because a face is constructed or put together by combining or adding together the (weighted) eigenvectors. Two aspects of this approach differ in useful ways from more traditional ideas of feature representation of faces. First, these eigenvectors are dependent on the statistical properties of the set of faces used to reconstruct the covariance matrix. This is an important characteristic of human face recognition. For example, it is well known that people are better able to recognize faces of their own race than faces of other races. We have modeled this recently as a problem in perceptual learning by using a principal components approach.⁷ By this account, faces of different races can be thought of as different statistical categories of faces (i.e., different races of faces may differ in norms for, e.g., face shape and eye/hair color). When a covariance matrix of faces is constructed by using a majority of faces of one race and a minority of faces of another race, the eigenvector representation, which is dependent on the statistical characteristics of the stimulus set, will capture subtle variations in the form and configuration of the majority-race faces. However, this representation will be less able to capture featural information that is useful for discriminating among minority-race faces. Thus, according to this model, the difficulties that we have in recognizing other-race faces arise from the inappropriate use of a set of features that are derived from, and hence optimized for, discriminating among faces of one's own race. A principal components approach models this general process in a relatively natural way.

A second advantage of using eigenvectors to represent faces is that, because they span the entire face, they can capture more global aspects of the faces than can local feature-based descriptions. Models that use

representations of local features must generally add in an explicit mechanism to represent configurational information in faces.

We demonstrate that, within the principal components framework, different bands of eigenvectors are differentially useful for recognition and categorization along visually derived semantic dimensions. In particular, a low-dimensional representation of faces in the lower dimensions of the space^{9,10} (i.e., a representation that use a subset of eigenvectors with the largest eigenvalues) has been proposed as a way of representing face images optimally in the least-squares error sense. Whereas this representation minimizes the physical similarity (squared) error between the original and low-dimensional representation, we show that it is not necessarily the most efficient representation for the task of recognition. In some cases, a low-dimensional representation of faces in the higher dimensions of the space (i.e., the dimensions associated with the eigenvectors with smaller eigenvalues) can give equally good, if not better, information for face recognition. We shall likewise provide a simulation and a visual demonstration, consistent with earlier work⁸, of the importance of the lower-dimensional eigenvectors for the prediction of general physical categories such as the sex of a face.

2. PRINCIPAL COMPONENTS APPROACH

The principal components approach to representing faces has been presented in detail elsewhere^{6,9,10} so we present only a short history and summary here. What is frequently forgotten in the current literature is that the approach is completely consistent with much older work by Kohonen¹² on autoassociative memory matrices, which he applied to the storage and retrieval of face images. Briefly, in this conceptualization, a digitized image of each face is coded as a vector that comprises pixel elements concatenated from the rows of the face image. Thus, the i th face is represented by a $J \times 1$ vector (where J is equal to the width times the height of the face image in pixels) and is denoted by \mathbf{a}_i . For convenience, normalized vectors are assumed (i.e., $\mathbf{a}_i^T \mathbf{a}_i = 1$). From the face vectors, an autoassociative matrix is constructed as follows:

$$(1) \quad \mathbf{W} = \sum_i \mathbf{a}_i \mathbf{a}_i^T .$$

Recall (i.e., reconstruction) of a face from the matrix is achieved as follows:

$$(2) \quad \hat{\mathbf{a}}_i = \mathbf{W} \mathbf{a}_i$$

where $\hat{\mathbf{a}}_i$ is the estimate of \mathbf{a}_i . The quality of this estimate is evaluated by computing the cosine between the original and reconstructed image vectors; formally:

$$(3) \quad \cos(\hat{\mathbf{a}}_i, \mathbf{a}_i) = \frac{\hat{\mathbf{a}}_i^T \mathbf{a}_i}{\|\hat{\mathbf{a}}_i\| \|\mathbf{a}_i\|} = \frac{\hat{\mathbf{a}}_i^T \mathbf{a}_i}{\|\hat{\mathbf{a}}_i\|} .$$

Because each term $\mathbf{a}_i \mathbf{a}_i^T$ of Eq. (1) is an outer product, the matrix \mathbf{W} is semipositive definite. The eigendecomposition of this autoassociative matrix which is equivalent to principal components analysis⁴, can be expressed as:

$$(4) \quad \mathbf{W} = \sum_j^r \lambda_j \mathbf{p}_j \mathbf{p}_j^T$$

where \mathbf{p}_j is the j th eigenvector of \mathbf{W} , λ_j is the eigenvalue associated with the j th eigenvector, and r is the rank of \mathbf{W} . Hence, Eq. (2) can be rewritten as

$$(5) \quad \hat{\mathbf{a}}_i = \sum_j^r \lambda_j \times (\mathbf{a}_i^T \mathbf{p}_j) \times \mathbf{p}_j .$$

Therefore, faces can be represented as a set of weights or coefficients on the set of eigenvectors. We refer to the set of coefficients that describes a particular face as its coefficient profile⁹.

The storage capacity of this matrix is approximately 15% of its dimensionality¹³. This estimate is for random vectors, however, and faces are highly correlated. The correlated nature of the face images decreases the capacity estimate. On the other hand, the Widrow-Hoff error-correction rule (or delta rule, as it is sometimes called) is frequently applied iteratively to optimize the quality of recall across the stimulus set. This increases the capacity of the matrix. In Eq. (5) this is equivalent to dropping the eigenvalues when reconstructing the images. In any case, because the dimension of our images, and nearly all of the face stimuli used in related work, are very large compared with the number of stimuli, this limit is not generally a problem.

3. HISTORY OF THE APPLICATION OF THE APPROACH

Kohonen¹² demonstrated that the autoassociative memory could act as a content-addressable memory. He demonstrated this by storing face images in the

memory and by reconstructing and displaying pixel vector reconstructions when noisy or incomplete face vectors were given as memory keys. In general, these reconstructions are convincingly similar to the original stored pixel vectors. Formally, this amounts to cosines between the original and reconstructed image vectors that are generally close to 1.0.

3.1. Recognition

Kohonen, however, did not demonstrate that the autoassociative matrix can act as a system for recognizing faces. To do so, it is necessary to show that the model can distinguish between learned and novel faces. Because all faces are highly correlated, reconstruction of novel faces using Eq. (2) generally produces images that are clearly facelike and that generally resemble the unlearned original face. To demonstrate that the model is able to distinguish learned from new faces, the quality of the reconstructions for the learned faces should, on the average, exceed the quality of the reconstructions for the unlearned faces. This was shown recently in a study³ that compared cosines between original and reconstructed face images for both learned and novel faces. On the average, cosines for the learned faces exceeded those for the novel faces, indicating that the model, within capacity limits, could distinguish learned from novel faces by using these cosines. We will provide a detailed explanation of this recognition testing procedure in the context of the simulations' Method section below.

3.2. Low dimensional representation

Recently, Sirovich and Kirby¹⁰ proposed a low-dimensional representation of faces using a subset of the first n eigenvectors, where ($n < r$) and r is the rank of the matrix. This representation is optimal in

the least-squares sense and hence minimizes the error of the reconstructed images. Sirovich and Kirby demonstrate convincingly, both mathematically and visually, that only a subset of the terms of Eq. (5) are needed to produce a good likeness of the original learned face. In some cases no perceptible differences can be seen between reconstructions made with the use of a subset of eigenvectors and the original faces. Whereas this representation is optimal in the least-squares error sense, the results of the simulations that we present below show that this representation is perhaps not the most useful one for recognition. The research is motivated by an informal observation that we made with the use of an XWindow-based tool for constructing faces from the eigenvector¹⁴. We noted that, in eliminating eigenvectors from the reconstructed faces, the likeness or general quality of the face decreased more by eliminating some ranges of higher eigenvectors (i.e., those with smaller eigenvalues) than by eliminating the more important lower eigenvectors (i.e., those with larger eigenvalues). For example, Fig. 1(a) displays an original face in the left panel, a reconstruction eliminating the first 20 eigenvectors in the center panel, and a second reconstruction eliminating the first 40 eigenvectors in the right panel. The identity of the face is barely compromised in these two reconstructions. On the other hand, Fig. 1(b) displays the original in the left panel, the face displayed with the first 20 eigenvectors in the center panel, and the face displayed with the first 40 eigenvectors in the right panel. Most human observers would have difficulty identifying the faces in these two reconstructions. This observation led us to examine the importance of different ranges of eigenvectors for recognition.

4. SIMULATION 1

4.1. Stimuli

One hundred fifty-nine faces were digitized from slides with a resolution of 16 gray levels. The faces were those of young Caucasian adults; roughly half were male and half were female. None of the slides presented people with facial hair or glasses. The images were aligned so that the eyes were at approximately the same height. The images were cropped around the face to eliminate clothing. Each face was

151 pixels wide and 225 pixels long, and so was represented by a 33,975-pixel vector consisting of the concatenation of the pixel rows. The simulations were carried out on a Sun Microsystems SparcStation and on a Convex C-1 Vector computer.

4.2. Method

An autoassociative memory was constructed with 100 randomly chosen faces (50 male and 50 female) from



(a)



(b)

FIGURE 1. (a) Left panel contains the original face. The center panel contains the face reconstructed by eliminating the first 20 eigenvectors. The right panel contains the face reconstructed by eliminating the first 40 eigenvectors. (b) Left panel contains the original face. The center panel contains the face reconstructed with the first 20 eigenvectors. The right panel contains the face reconstructed with the use of the first 40 eigenvectors

the original set. The remaining 59 pictures were reserved as novel faces for testing purposes. The autoassociative matrix was decomposed into its eigenvectors. In each part of this simulation, all 159 faces were reconstructed with Eq. (5) by using a range of eigenvectors. These ranges were as follows: a.) eigenvectors 1-15; b.) eigenvectors 5-20; c.) eigenvectors

10-25; and so on, up to and including r.) eigenvectors 85-100. Finally, a control simulation was carried out with the use of the entire set of eigenvectors.

The matrix was tested for accuracy using signal detection methodology¹⁵. Essentially, this procedure defines the signal as the old or learned faces and the noise as new or unlearned faces. By setting a quality

Eigenvector Range	d'	Hit Rate	False Alarm Rate	Percent Correct	Mean Cosine Old	Mean Cosine New
1-15	1.69	0.91	0.35	81.1	0.97	0.95
5-20	0.95	0.62	0.25	66.6	0.23	0.19
10-25	0.88	0.62	0.27	66.1	0.18	0.15
15-30	1.26	0.63	0.16	70.4	0.16	0.12
20-35	1.42	0.69	0.17	74.2	0.15	0.11
25-40	1.41	0.75	0.22	76.1	0.13	0.09
30-45	1.69	0.81	0.22	79.8	0.12	0.08
35-50	1.62	0.78	0.19	79.2	0.11	0.07
40-55	1.67	0.76	0.15	79.2	0.10	0.07
45-60	1.88	0.79	0.14	81.7	0.10	0.06
50-65	2.04	0.77	0.08	82.3	0.09	0.06
55-70	2.03	0.79	0.10	83.0	0.09	0.06
60-75	1.76	0.79	0.17	80.5	0.08	0.05
65-80	1.92	0.80	0.14	82.3	0.08	0.05
70-85	1.60	0.73	0.15	76.0	0.07	0.05
75-90	1.64	0.69	0.12	76.1	0.07	0.05
80-95	1.32	0.72	0.22	74.5	0.07	0.05
85-100	1.00	0.72	0.22	68.5	0.06	0.05
1-100	> 4.00	0.99	0.01	98.7	1.00	0.97

TABLE 1. Recognition as a Function of Eigenvector Range

of reconstruction (or cosine) criterion, the procedure divides the faces into old and new faces by assigning faces for which the quality of reconstruction measure exceeds the criterion to the old face group, and assigning those less than the criterion to the group of new faces. The procedure was implemented as follows: 1.) All faces were reconstructed in all of the eigenvector ranges noted above by using Eq. (5). 2.) The quality of the reconstructions was measured as the cosine between the original and reconstructed images [Eq. (3)]. 3.) The mean of the cosines between original and reconstructed faces for the learned and novel faces was calculated; as was the mean cosine between original and reconstructed faces. The mean of

these two means was used as the criterion or threshold for deciding whether a face was to be classified as old or new. The distributions of cosines for old and new faces were approximately normal, making this a sensible criterion. 4.) Model responses were categorized as hits (i.e., old faces that the model assigned to the old category) and false alarms (i.e., new faces that the model assigned to the old category). 5.) d' was calculated as $z(\text{hit rate}) - z(\text{false alarm rate})$. This amounts, simply, to the distance in z -score units between the means of the old and new cosine distributions. High d' 's indicate little overlap between the distributions for old and new faces, and hence high discriminability, whereas low d' 's indicate much overlap and, hence, poor discriminability.

4.3. Results

Table 1 contains the d' , hit and false alarm rates, and the overall percent correct classifications, as well as the mean cosines between original and reconstructed images for the old and new faces. Several things are worth noting. First, as might be expected from

earlier work¹⁰, the average physical similarity of the original face to the reconstructed face (as measured by the cosine between the original and reconstructed faces) decreases as the range of eigenvectors used shifts from the eigenvectors with larger eigenvalues

to those with smaller eigenvalues. The reconstructions of both the old and new faces at the top of this range (1 – 15) were nearly perfect. Reconstructions at the lower part of the eigenvector range, however, were very dissimilar to the original faces with average cosines only slightly above zero. Figure 2 shows the falloff of the quality of the reconstructions as a function of the eigenvector range.

The second point to note is that the quality of the reconstructions was consistently higher for the old faces than for the new faces, across the eigenvector ranges. Error bars are not included in Fig. 2 because even with the lowest of the d' s in Table 1 the model's ability to discriminate old from new faces is significantly greater than chance (Binomial test, $N = 159$, $p < .05$).

The third and most important point to notice from Table 1 is that despite this general monotonic increase in error of reconstruction as the range of eigenvectors changes to include only those in the higher dimensions of the face space, the ability of the model to

discriminate between old and new faces does not similarly decrease. There is a dropoff in discriminability in the second two eigenvector ranges and then a steady increase in discriminability that peaks in the 50-70 range. The pattern of d' values is illustrated in Fig. 3.

4.4. Discussion

While the reconstructions of the faces were physically more similar to the originals for the lowest range of eigenvectors, this information was not necessarily the most efficient for discriminating the learned from unlearned faces. In general, any of the 15-eigenvector ranges between 45 and 80 provided better information for discriminating learned from unlearned faces than did the first 15-eigenvector range. Intuitively, this makes sense because the eigenvectors with the larger eigenvalues are likely to convey information that is common to all the faces, whereas the eigenvectors with the smaller eigenvalues are likely to convey information about individual faces. This information about individual faces is likely to be the most useful for recognition.

5. SIMULATION 2

It has been shown⁸ that certain eigenvectors contain information relevant to physical categories of faces like race and sex. Using an autoassociative network with Caucasian and Japanese faces, O'Toole *et al.*⁸ found that the first, second, and seventh eigenvectors contained most of the information useful for predicting the race of the face. In particular, the mean coefficients for these three eigenvectors differed significantly for Caucasian and Japanese faces. By using only the coefficients for the second eigenvector, along with a threshold rule, the race membership of 167 Caucasian and Japanese faces was predicted with 88.6% percent accuracy. This study, however, looked only at the usefulness of the first seven eigenvectors in predicting race membership. The present simulation examined the full range of eigenvectors and provides a demonstration of the importance of a relatively small number of eigenvectors (those with relatively large eigenvalues) in controlling the sex of the reconstructed faces.

5.1. Method and Results

The usefulness of each eigenvector coefficient for discriminating faces based on sex was examined. There are a variety of methods for accomplishing this. One of the simplest ways is to use the value of the coefficients or weights of each of the faces on each of the eigenvectors for predicting the sex of the faces. Assigning zeros to female faces and ones to male faces, we computed the point biserial correlation coefficient, r , and the proportion of explained variance r^2 between sex and the eigenvector weight for each of the 159 faces. This is equivalent to performing a t -test using sex as an independent variable and the weights as a dependent variable. In other words, we wished to examine the degree of relationship between the inclusion of individual eigenvectors in the reconstruction and the sex of the face. Figure 4 shows the cumulative proportion of explained variance (i.e., r^2) for sex classification across the eigenvectors¹⁶. Only eigenvector coefficients that correlated significantly (significance of r test, $p < .05$) are included. Of the 100 eigenvectors, 12 showed a significant correlation between

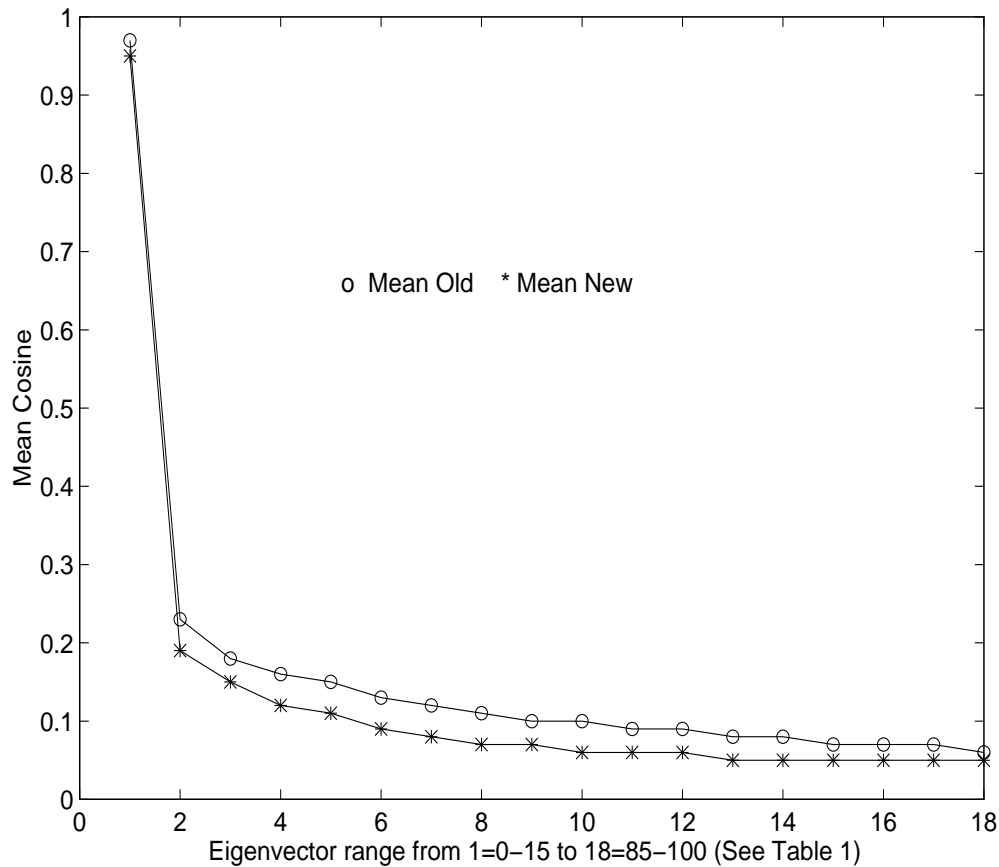


FIGURE 2. Mean quality of reconstructions, as measured by cosine between the original and reconstructed images, for the old and new faces as a function of eigenvector range. As the range moves from eigenvectors with larger to smaller eigenvalues, the quality of reconstructions falls off exponentially. However, the quality of reconstruction for old faces exceeds that for new faces in all eigenvector ranges, indicating an ability to discriminate between learned and new faces by using the quality of reconstruction.

sex and eigenvector weight. Of these, 5 appear in eigenvectors with the 15 highest eigenvalues. Additionally, as can be seen in Fig. 4, the largest amount of explained variance of sex prediction occurs in the eigenvectors with the largest eigenvalues. In particular, the second eigenvector was the best predictor of the sex of the face ($r = .66$; $r^2 = .38$). A striking demonstration of the importance of this eigenvector

in determining the sex of the face appears in Fig. 5. Figures 5(a) and 5(b) contain the first and second eigenvectors, respectively. Whereas the second eigenvector appears monsterlike, adding it to the first, as is displayed in the Fig. 5(c), creates a masculine-looking face. However, subtracting the second eigenvector from the first, as is displayed in Fig 5(d) produces a feminine looking face.


Mean d' as a Function of Eigenvector Range

FIGURE 3. Recognition ability of the model as measured by d' . Recognition performance drops off initially as the range moves from eigenvectors with larger to smaller eigenvalues but rebounds in the higher ranges, exceeding the initial d' value.

6. GENERAL DISCUSSION AND CONCLUSION

This study shows that in choosing an optimal representation for faces within the principal components approach, the particular task to be performed must be taken into account. The base physical similarity of the represented images to the originals does not provide the best measure of useful information for distinguishing faces from one another. A low-dimensional representation of the faces in the dimensions of the space associated with small eigenvalues provides good information for the identification of specific faces, for discriminating between learned and novel faces, and also for the human judgment of similarity between the reconstructed face and the target. It is interesting to compare these findings with some other theoretical models of face processing. For example, Bruce

and Young¹¹, in their often-cited model of face processing, distinguish between identity-specific semantic codes and visually-derived semantic codes. They draw the distinction in a box and arrows framework. This distinction seems to apply directly to the present model. Clearly, the network model spontaneously implements the dissociation between the two varieties of visual information in faces to which identity-specific and facial-category labels would be appropriately associated. Therefore, the model can be seen as more parsimonious than a box and arrow model.

One important implication of this paper is to stress the importance of the task at hand when discussing the optimal properties of a model. The least-squares criterion seems particularly well adapted when the problem is to reconstruct a face, or to derive some




FIGURE 4. Cumulative proportion of explained variance for sex prediction as measured by r^2 between sex and eigenvector weight for each eigenvector. Weights on the eigenvectors with the largest eigenvalues account for most of the explained variance.

type of common semantic information (e.g., sex, race). However, the same criterion leads to ignoring the identity-specific information required to minimize confusions between new and old faces when the task is to recognize faces.

REFERENCES

- [1] L. D. Harmon, "Some aspects of recognition of human faces," in *Pattern Recognition in Biological and Technical Systems*, O. J. Grusser and R. Klinke, eds. (Springer-Verlag, Berlin, 1971),
- [2] M. A. Shackleton and W. J. Welsh, "Classification of facial features for recognition," in *Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 1991), pp. 573-579.
- [3] A. J. O'Toole, R. B. Millward, and J. A. Anderson, "A physical system approach to recognition memory for spatially transformed faces," *Neural Networks*, **1**, 179-199 (1988).
- [4] H. Abdi, "A generalized approach for connectionist auto-associative memories: interpretation, implications and illustration for face processing," in *Artificial Intelligence and Cognitive Sciences*, J. Demongeot, ed. (Manchester University Press, Manchester, 1988).
- [5] B. A. Golomb, D. T. Lawrence, and T. J. Sejnowski, "SEXnet: A neural network identifies sex from human faces," in *Advances in Neural Information Processing Systems*, **3**, D. S. Touretsky and R. Lippmann, eds. (Morgan Kaufmann: San Mateo, CA, 1991).
- [6] R. B. Millward and A. J. O'Toole, "Recognition memory transfer between spatially transformed faces," in *Aspects of Face Processing*, H. D. Ellis, M. A., Jeeves, F. Young, and A. Newcombe, eds. (Martinus Nijhoff: Dordrecht, 1986).




FIGURE 5. (a) First eigenvector. (b) Second eigenvector which appears monsterlike. (c) First eigenvector plus the second; the resulting face appears masculine. (d) First eigenvector minus the second, the resulting face appears feminine.

- [7] A. J. O'Toole, K. Deffenbacher, H. Abdi, and J. C. Bartlett, "Simulating the 'other-race effect' as a problem in perceptual learning," *Connection Science*, **3**, 163-178 (1991).
- [8] A. J. O'Toole, H. Abdi, K. A. Deffenbacher, and J. C. Bartlett, "Classifying faces by race and sex using an autoassociative memory trained for recognition," in *Proc. 13th Annual Conf. of the Cognitive Science Society* (Chicago, Illinois, 1991).
- [9] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cognitive Neuroscience*, **3**, 71-86 (1991).
- [10] L. Sirovich and M. Kirby, "Low-dimensional procedure for the characterization of human faces," *J. Opt. Soc. Am.*, **4**, 519-518 (1987).
- [11] V. Bruce and A. W. Young, "Understanding face recognition," *British Journal of Psychology*, **77**, 305-327 (1986).
- [12] T. Kohonen, *Associative Memory: A System Theoretic Approach* (Springer-Verlag, Berlin), 1977.
- [13] J. J. Hopfield, "Neurons with graded responses have collective computational properties like those of two-state neurons," *Proceedings of the National Academy of Sciences*, **81**, 3088-3092 (1984).
- [14] A. J. O'Toole and J. Thompson, "An X Windows-Based tool for synthesizing face image from eigenvectors," *Behavior Research Instruments and Computers* (in press, 1992).
- [15] D. M. Green and J. A. Swets, *Signal Detection Theory and Psychophysics*. (John Wiley : New York), 1966.
- [16] This display is used because multiple significance tests that employ a set probability increase the possibility of rejecting the null hypothesis by chance across the string of comparisons.