Connection Science, 1991, **3(2)**, 163–178.

Simulating the "other-race effect" as a problem in perceptual learning

Alice O'Toole
The University of Texas at Dallas

Ken Deffenbacher
The University of Nebraska at Omaha

Hervé Abdi
The University of Texas at Dallas

Jim Bartlett
The University of Texas at Dallas

We report a series of simulations on the well-known "other-race effect." We trained an autoassociative network on a majority and a minority race of faces, and tested the model's ability to process faces from the two races in different ways. First, the model was better able to reconstruct unlearned majority faces than minority faces. Secondly, the average inter-face similarity was higher for the reconstructed minority faces than for reconstructed majority faces, indicating that the model was coding the majority faces more distinctively than the minority faces. These results held for Caucasian faces as the majority race and Japanese faces as the minority race and vice versa. Thirdly, we simulated a recognition task for same- and other-race faces by using a face history matrix and a recognition task matrix with equal numbers of Caucasian and Japanese faces, and reconstructing these faces as a weighted combination of the two matrices. Using Caucasian faces as the majority race, the model was better able to discriminate learned from new Caucasian faces than learned from new Japanese faces. We discuss the results in terms of perceptual tuning to information useful for processing faces of a single race.

Keywords: Face memory, autoassociative memory, neural network, other-race effect. .

Introduction

For many years scientist and layperson alike have suspected that faces of one's own race are recognized more accurately than faces of another race. In the early part of the century, Feingold (1914,p. 50) stated the supposition and a plausible reason for its existence this way

> other things being equal, individuals of a given race are distinguishable from each other in proportion to our familiarity, to our contact with the race as a whole. Thus, to the uninitiated American all Asiatics look alike, while to the Asiatic, all White men look alike.

More recently, it has been shown that approximately half of potential jurors believe that such a bias exists (Deffenbacher & Loftus, 1982). Indeed, abundant empirical support for the own-race bias in face recognition accuracy can be found in the recent meta-analyses of a large number of studies

Thanks are due to June Chance and Al Goldstein for providing the Caucasian and Japanese faces used in the simulations, and to Peter Assmann, Barbara Edwards and two anonymous reviewers for helpful comments on an earlier version of this manuscript. This project was supported in part by National Institute of Aging Grant RO1-AG07798 to J.C.B. Send correspondence about this paper to Hervé Abdi: `herve@utdallas.edu`. Home page: `http://www.utdallas.edu/∼herve`.

on the topic (Shapiro & Period, 1986; Bothwell et al., 1989).

Several hypotheses have been advanced to account for the cross-race phenomenon: Faces of some races are inherently more difficult to identify than others; prejudicial attitudes lead to less accurate recognition for other-race faces; and other-race faces are processed more superficially than same-race faces. There is little support for any of these accounts (Brigham, 1986). A fourth possibility is that implied by the quote from Feingold (1914): an own-race bias is in direct proportion to the difference in amount of contact with persons of one's own and another race. Several studies have found a smaller other-race effect for persons living in more racially integrated circumstances (Cross et al., 1971; Feinman & Entwistle, 1976). However, while at least one study (Brigham et al., 1982) yielded a small but significant correlation of self-reported degree of cross-race contact and cross-race recognition ability, other studies (e.g. Brigham & Barkowitz, 1978) have found no relationship at all.

It may be, however, that current techniques are not sensitive enough to adequately assess the quantity and quality of contact with persons of another race. There are data which suggest that the cross-race effect may indeed be a matter of differential exposure to faces of different races. For one thing, a number of attempts to improve same-race face recognition have all failed (Malpass, 1981). However, similar training efforts for other-race face recognition have yielded improvements (e.g. Goldstein & Chance, 1985).

We think that the differential effects of further training in face recognition are due to differential amounts of perceptual learning associated with same versus other-race faces. A cross-sectional study of the development of face recognition ability by Chance et al. (1982) found that 6 year old Caucasian children show only a small cross-race effect recognizing Japanese compared with Caucasian faces. At successively older ages up to early adulthood, ability to recognize both races increased, but ability to recognize same-race (Caucasian) faces increased much more rapidly. Hence, any attempt to improve same-race face recognition by short-term training programs may be inadequate compared with years of extensive processing of same-race faces.

Studies examining the role of perceptual learning[1] in the other-race effect are difficult to carry out empirically for two reasons. First, while it may be possible to find subject populations with a relatively controlled "face-learning" history, it is generally not possible to equate the populations along other important cultural and social dimensions that may affect performance on the task. Secondly, as we have already noted, short-term perceptual learning studies involving practice with a single race of faces are not necessarily adequate to control for the lifetime experience of observers with faces of their own race. These methodological difficulties make the cross-race effect an ideal candidate for simulation approaches to understanding the psychological data.

In the present study, we present simulations of a perceptual learning account of the other-race effect that is based on the following principles. First, we assume that faces of different races comprise different statistical categories of faces. Secondly, within a given category of faces, a set of differentially weighted "features[2]" is optimal for encoding faces in a manner that makes faces within the category most discriminable. Different feature sets and weightings, however, are optimal for processing faces from other-race categories of faces. Thirdly, with exposure to many faces of a given race and a smaller number of faces of other races, perceptual learning enables observers to make optimal use of the features that are best for processing faces from the category with which they have had the most experience, topically, faces of their own race. By this account, the difficulties experienced with faces of another race are due to the fact that the optimal features for distinguishing faces of one's own race are not optimal in processing the faces of another race.

One way to simulate the other-race effect is to train an autoassociative network on different proportions of faces of an 'own' and 'other' race. We

[1] Referred to in the face recognition literature simply as "experience."

[2] The word features is used in its most general sense without commitment to a specific definition.

trained an autoassociative system on a large number of faces of a majority race and a smaller number of faces of a minority race to mimic the other-race effect. The advantages of an autoassociative memory used with Widrow-Hoff error correction is that it will develop connection weights in such a way as to optimize the storage capacity of the matrix. Thus, with very similar stimuli, such as a single race of faces, the model should tune itself to the information important for processing faces from within the class. Due to the distributed nature of the memory, when faces are retrieved from the system, they will be filtered by the learning history of the system.

Several predictions about the system's ability to process same- and other-race faces follow. First, when the network is trained on a majority of faces of one race and a minority of faces of another race, its ability to represent faces of the majority race should be better than its ability to represent faces from the minority race. This is due to the fact that model will have developed "features" that are more appropriate for faces of the majority race. We can assess the validity of this prediction by looking at the quality of face reconstructions for new (previously unencountered) faces of the majority and minority races. Secondly, reconstructed new faces of the minority race should be more similar to one another than reconstructed faces of the majority race. In other words, the average inter-face similarity should be greater for reconstructed minority faces than for reconstructed majority faces. This is because the model will not develop a coding that makes optimal use of the distinguishing features for the minority race; hence, these faces should be "perceptually" more similar to one another. Finally, the model should be better able to recognize majority faces than minority faces.

The simulations serve, first, to test the model qualitatively as a face recognition tool with a much larger and higher quality stimulus set than that used previously (Kohonen, 1984; O'Toole et al., 1988; O'Toole & Abdi, 1989). We will look specifically at the model's performance with respect to the predictions stated above.

Secondly, this type of model suggests a different definition of features than has previously been used to characterize faces. Since the autoassociative memory can be decomposed into a set of eigenvectors, and since faces learned by the model can be reconstructed by the weighted combination of these eigenvectors, the eigenvectors may be thought of as features for characterizing the stimulus set. We should expect to see differences in eigenvectors based on the face history of the model. Furthermore, since we used a simple visual code in these simulations, the eigenvectors can be displayed as images. We shall discuss the potential role of the eigenvectors as features for characterizing same-and other-race faces.

Simulation 1

The model is defined first and then its application to the other-race problem is presented. A digitized image of each face was coded as a vector comprised of pixel elements concatenated from the rows of the face image. Thus, the ith face was 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 $\mathbf{f}_i$. For convenience, normalized vectors are assumed (i.e. $\mathbf{f}_i^T \mathbf{f}_i = 1$). The autoassociative matrix was constructed as

$$\mathbf{A} = \sum_i \mathbf{f}_i \mathbf{f}_i^T \tag{1}$$

Recall of individual faces from the matrix was done according to the rule

$$\widehat{\mathbf{f}_i} = \mathbf{A}\mathbf{f}_i \tag{2}$$

where $\widehat{\mathbf{f}_i}$ is the system estimate of $\mathbf{f}_i$. The quality of this estimate is measured by comparing the reconstructed image with the original image using the cosine of the angle between the vectors $\widehat{\mathbf{f}_i}$ and $\mathbf{f}_i$. The Widrow-Hoff error-correction rule was applied iteratively to optimize the quality of the recall across the stimulus set

$$\mathbf{A}_{[t+1]} = \mathbf{A}_{[t]} - \gamma\left(\mathbf{f}_i - \mathbf{A}_{[t]}\mathbf{f}_i\right)\mathbf{f}_i^T \tag{3}$$

where i is randomly chosen and γ decreases as the reciprocal of the iteration number.

Since the eigen-decomposition of the autoassociative matrix is equivalent to principal component analysis (Abdi, 1988), the autoassociative matrix

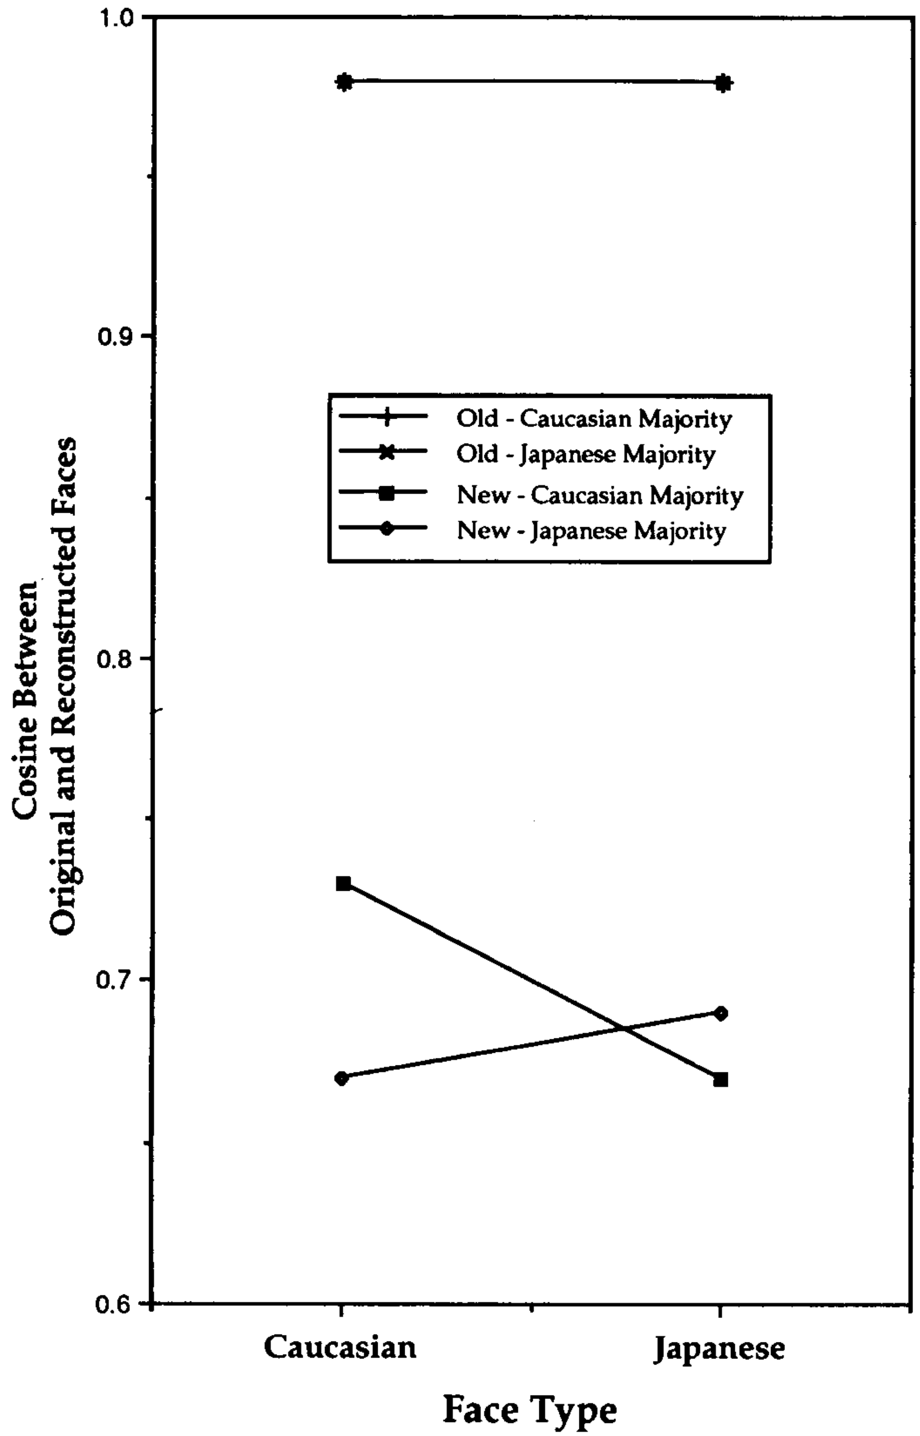

Figure 1. Mean cosines between original and reconstructed images for the OLD and NEW majority and minority faces.

can give indications of the statistical structure of the stimulus set. The storage capacity of such a matrix is approximately 15% of its dimensionality for random vectors (Hopfield, 1984). Since the dimension of our images was very large by comparison with the number of stimuli, this limit was not a problem for these simulations.

Method

Stimuli. A total of 319 Caucasian and Japanese faces were digitized from slides with a resolution of 16 grey levels using a Fotovix digitizer attached to a 286-based computer with a 16-bit TARGA board (True Vision). Faces were of young adults and were roughly half male and half female. None of the slides pictured people with facial hair or glasses. The images were aligned so that the eyes were at about 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 pixels rows. A spatial differentiation encoding was used to enhance lines prior to the extraction of the pixel vector (cf. O'Toole et al., 1988). The simulations were carried out on a Sun MicroSystems SparcStation and on a Convex C-1 Vector computer.

Procedure. Two simulations of the other-race effect were performed: one used Caucasian faces as the majority race and Japanese faces as the minority race, and the other used Japanese faces as the majority and Caucasian faces as the minority group. For the Japanese minority simulation, an associative memory was trained using error correction on 95 Caucasian and five Japanese faces. For the Caucasian minority simulation 95 Japanese and five Caucasian faces served as the training set[3].

Results and Discussion

Representations of majority and minority race faces. The model was tested by reconstructing the Japanese and Caucasian faces that the model learned (OLD), and by reconstructing a sample of Japanese and Caucasian faces not learned by the model (NEW). The cosine between the original and reconstructed image indicates the quality of the model's representation of the face. Figure 1 shows the mean cosines for the OLD and NEW majority and minority faces for the simulations. Three points are worth noting. First, in both simulations, the OLD stimuli (both Caucasian and Japanese faces) were nearly perfectly reconstructed (mean cosine=0.98). This is a consequence of the fact that the capacity of the matrix was not challenged (cf. Hopfield, 1984, and below). We discuss below one method of degrading the performance of the model in a psychologically interesting way. Secondly, the average cosine for the reconstructed NEW majority faces was greater than the average cosine for the reconstructed NEW minority faces. This can be seen in the interaction in Figure 1. The differences between the quality of the reconstructions for majority and minority faces reflects the model's greater success in coding or representing novel faces from the majority race than from the minority race.

Finally, in both simulations, the cosines for the NEW faces did not reflect random performance for the model. In other words, the minority race faces were not completely unfamiliar stimuli for the model. This is a consequence of the fact that all faces share a general schema of features and so a given race might be best thought of as a subcategory of the general class of face stimuli.

Similarity. We tested the prediction that novel majority faces are perceived by the model to be less similar to one another (i.e. more distinctive) than minority faces. An analysis of the similarity of the reconstructed NEW faces to one another was carried out. In this analysis, 50 randomly chosen NEW Caucasian faces and 50 randomly chosen NEW Japanese faces were reconstructed and the model's estimate of each face was used for this similarity analysis. These stimuli can be thought of as filtered or "perceived" by the matrix trained with a majority and minority race. For each race of faces, the inter-face similarity for the recalled faces

[3] The choice of 95% majority and 5% minority faces is arbitrary. We have carried out the first set of simulations with 75% and 25%, as well, and have found qualitatively similar, though less extreme, results for the model's ability to represent new majority and minority race faces.

was calculated by taking the cosine between all possible pairs of different 'perceived' faces. The cosine between two faces indicates the similarity between the two, with identical or scaled faces yielding cosines of 1.0. The average similarity of all possible pairs of reconstructed faces was taken as the average inter-face similarity.

Several conditions were analyzed. For each majority simulation, the average inter-face similarity of Caucasian and Japanese faces was computed. Furthermore, the reconstructions were carried out using different numbers of eigenvectors to look at the consistency of the similarity effects. To explain this latter analysis, a short digression into the properties of associative matrices is necessary. The reconstruction of any face from the autoassociative matrix can be achieved either by Equation (2) or, equivalently, by taking a weighted sum of the eigenvectors of the matrix $\mathbf{A}$, where the weights of each eigenvector for a face $\mathbf{f}_i$ are equal to the dot-product between the face vector and the eigenvector (multiplied by the eigenvalue of the eigenvector—when error correction is not being used, since error correction has the effect of equalizing the eigenvalues). For these simulations, error correction was used and so the eigenvalue was not included in the weights. Thus, the reconstruction is given by

$$\widehat{\mathbf{f}_i} = (\mathbf{f}_i \cdot \mathbf{e}_1)\mathbf{e}_1 + (\mathbf{f}_i \cdot \mathbf{e}_2)\mathbf{e}_2 + \dots \\ \dots + (\mathbf{f}_i \cdot \mathbf{e}_\ell)\mathbf{e}_\ell + \cdots + (\mathbf{f}_i \cdot \mathbf{e}_n)\mathbf{e}_n \quad (4)$$

where $\mathbf{e}_\ell$ indicates the ℓ-th eigenvector.

The reconstruction of each face, then, can be quantified precisely by this list of coefficients and the set of eigenvectors of $\mathbf{A}$[4]. Returning from our digression, it is dear that a face may be recalled using this second procedure with all or any subset of eigenvectors. Since eigenvectors can be ordered by importance of contribution using their associated eigenvalues, we recalled faces using different numbers of eigenvectors to test the consistency of the results as more eigenvectors were included.

The results of this analysis appear in Figure 2 (a) for the Caucasian majority simulation and in Figure 2 (b) for the Japanese majority simulation. Average interface similarity, as defined by the average cosine between all possible pairs of reconstructed faces (as coded by the set of coefficients used to reconstruct them), appears on the y axis and number of eigenvectors is plotted on the x axis. In both simulations, faces from the minority race were more similar to one another on the average than were faces in the majority race. Thus, when the model is trained on a majority of faces of one race and a minority of faces of another race, it creates more distinct codings of majority race faces. This finding is reminiscent of Feingold's (1914) quote.

Simulation 2

As previously noted, the capacity of an autoassociative memory without error correction can be estimated as approximately 15% of its dimensionality. This estimate assumes random vectors. The vectors we used were of dimensionality 33 975 and so the capacity of the memory should be roughly 5096 faces. While there are two differences between these simulations and those for which the capacity estimates were derived, these have inverse effects on the capacity estimates. First, we used error correction, which improved the capacity of the matrix. Secondly, faces are not random vectors but are highly correlated, a factor that lessens the capacity. In any ease, it is clear from Simulation 1 that 100 faces did not challenge the capacities of the matrix. Since we had a limited data base of faces, we explored a number of methods for degrading the system's performance. At least one of these is interesting psychologically and would merit attention regardless of the performance constraints. This method draws on the metric multidimensional sealing analogy cited previously.

Multidimensional scaling tries to represent space relations between entities of a stimulus set in the smallest dimensional space possible while accounting for some experimenter-set criterion of variance. The eigenvectors of an associative mem-

[4] A strong analogy with metric multidimensional scaling is present here. The axes or dimensions of metric multidimensional scaling solutions are the eigenvectors ordered by the magnitude of their associated eigenvalues. Thus, the first axis is the first eigenvector, etc. Typically, multidimensional scaling solutions use as many axes as are needed to account for some experimenter-specified proportion of variance.

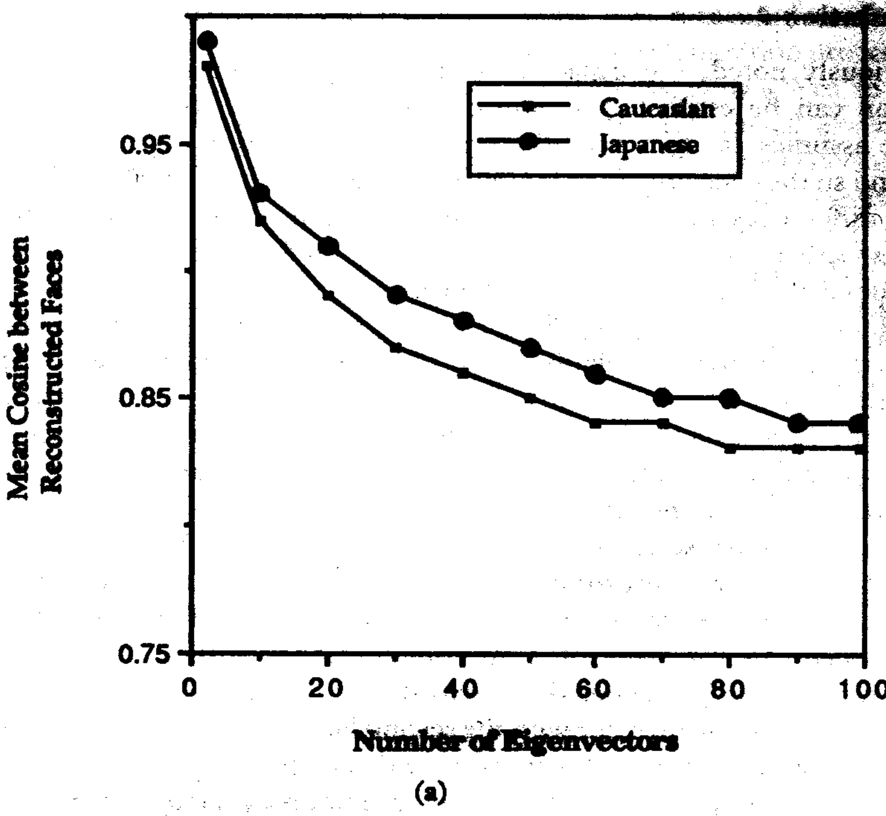

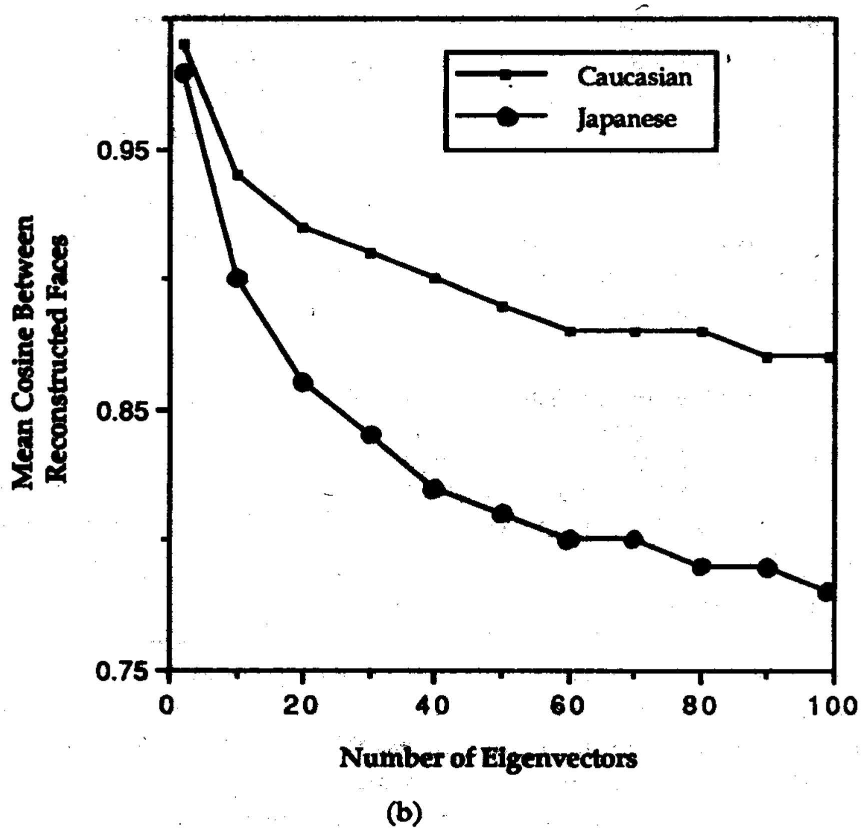

Figure 2. Average inter-face similarity for the (a) Caucasian and (b) Japanese majority (95%) simulations, plotted as a function of the number of eigenvectors used to reconstruct the faces. The minority (5%) race faces for both simulations are more similar to one another than are faces in the majority race.

ory are equivalent to the axes of a multidimensional scaling solution, with the eigenvector with the largest positive eigenvalue accounting for the largest proportion of variance, and the eigenvector with the second largest eigenvalue accounting for the second largest proportion of variance, and so on. The maximum number of dimensions needed to account for all of the variance is equal to the rank of the matrix (i.e. the number of eigenvectors with non-zero eigenvalues). Frequently in multidimensional scaling, however, an acceptably large proportion of the variance may be accounted for by a very small set of dimensions and, thus, the eigenvectors with smaller eigenvalues may be discarded without losing much information about the structure of the stimulus set.

Likewise, recall from an associative memory can be carried out using a smaller number of eigenvectors (cf. equation (4)). The criterion for an acceptable number of dimensions in this case, however, is one that maintains an acceptable (but not perfect) level of recognition performance.

Recognition Memory for Same- and Other-race Faces

To simulate a recognition memory task we need to model two components of memory, a long-term experience component (i.e. face race history) and a short-term face recognition task. We expect experience to affect the short-term recognition task in the ways outlined above. For the purpose of completeness, we report two simulations. In the first, we tested the ability of the autoassociative memory to distinguish between OLD and NEW faces for a majority matrix. In the second, we added a short-term component to this matrix, which consisted of half Caucasian and half Japanese faces. We then examined the ability of the model to discriminate OLD from NEW faces for these additional faces. We should note that we do not believe that this is the only or even best way to simulate such a task. We feel that it is the simplest way, however, and so we chose to explore this method first.

Method

The matrix was tested for accuracy with a Yes/No procedure as follows. Learned Caucasian and Japanese faces (OLD) and NEW Caucasian and Japanese faces were reconstructed. The quality of the reconstructions was measured as the cosine between the original and reconstructed images. A Yes/No recognition procedure was implemented by setting a criterion cosine value β and by assigning a 'Yes' to faces for which the cosine between the original and reconstructed image exceeded the criterion and 'No' to faces for which the cosine was less than the criterion β. The most direct choice for β is the mean of the cosine distribution means for the reconstructed OLD faces and the reconstructed NEW faces. Signal detection methodology maps easily onto this Yes/No task since the distribution of cosines for OLD faces can be thought of as the signal distribution and the distribution of cosines for NEW faces as the noise distribution. OLD faces with cosines greater than β are considered hits and NEW faces with cosines greater than β are considered false alarms. A d' score may then be computed in the standard way. Also, since the distribution of cosines for the signal (i.e. the OLD faces) and the distribution of cosines for the noise (i.e. the NEW faces) are known completely, a ROC curve may be plotted by choosing β values and calculating the hit and false alarm rates that would result from using these different criteria.

Results and Discussion

To test the accuracy of the models, we used all of the OLD faces (100 faces: 95 majority and five minority) and a sample of 120 NEW faces, approximately half Japanese and half Caucasian. The accuracy of the model using all the eigenvectors was essentially perfect. We degraded the simulations, therefore, by using smaller numbers of eigenvectors. Figure 3 (a and b) displays ROC curves for the performance of the Caucasian and Japanese majority models, respectively, with three different numbers of eigenvectors contributing to the reconstruction. For both majority simulations, 10 eigenvectors yielded excellent performance. Dividing the faces into majority and minority face groups did not show the cross-race effect. That is, majority faces did not yield larger values of d'. This is likely to be due to the fact that only five minority-

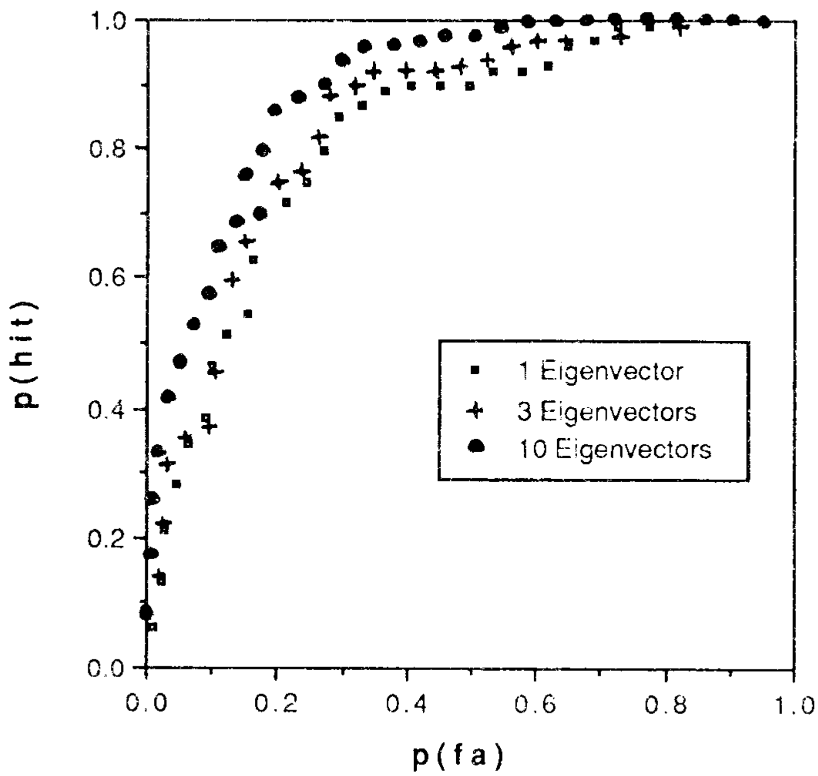

(a)

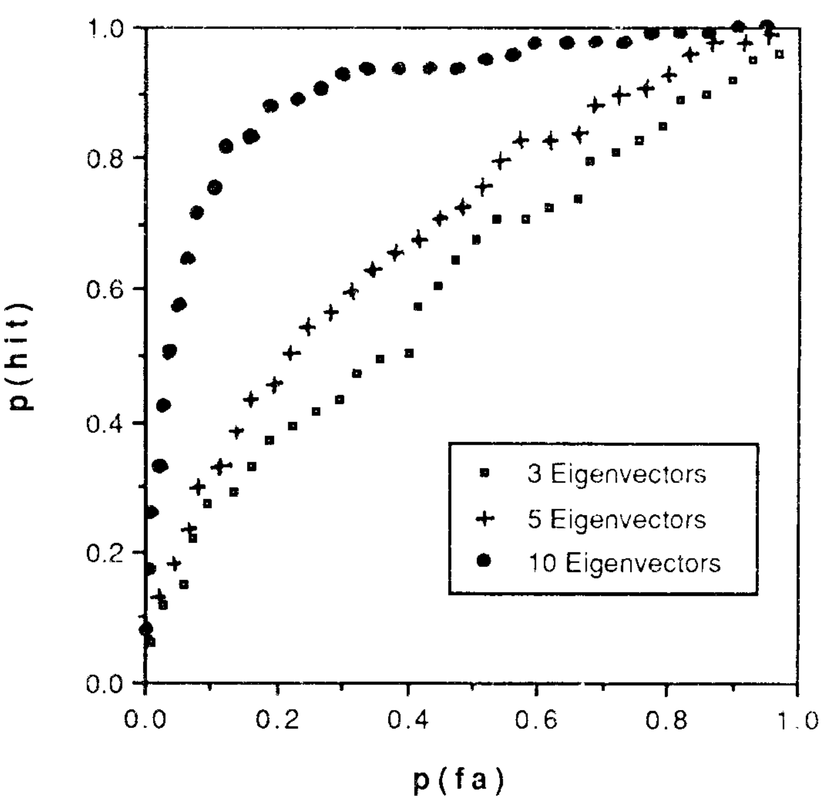

(b)

Figure 3. ROC curves for the performance of the (a) Caucasian and (b) Japanese majority (95%) models. Performance is plotted with different numbers of eigenvectors contributing to the reconstructions.

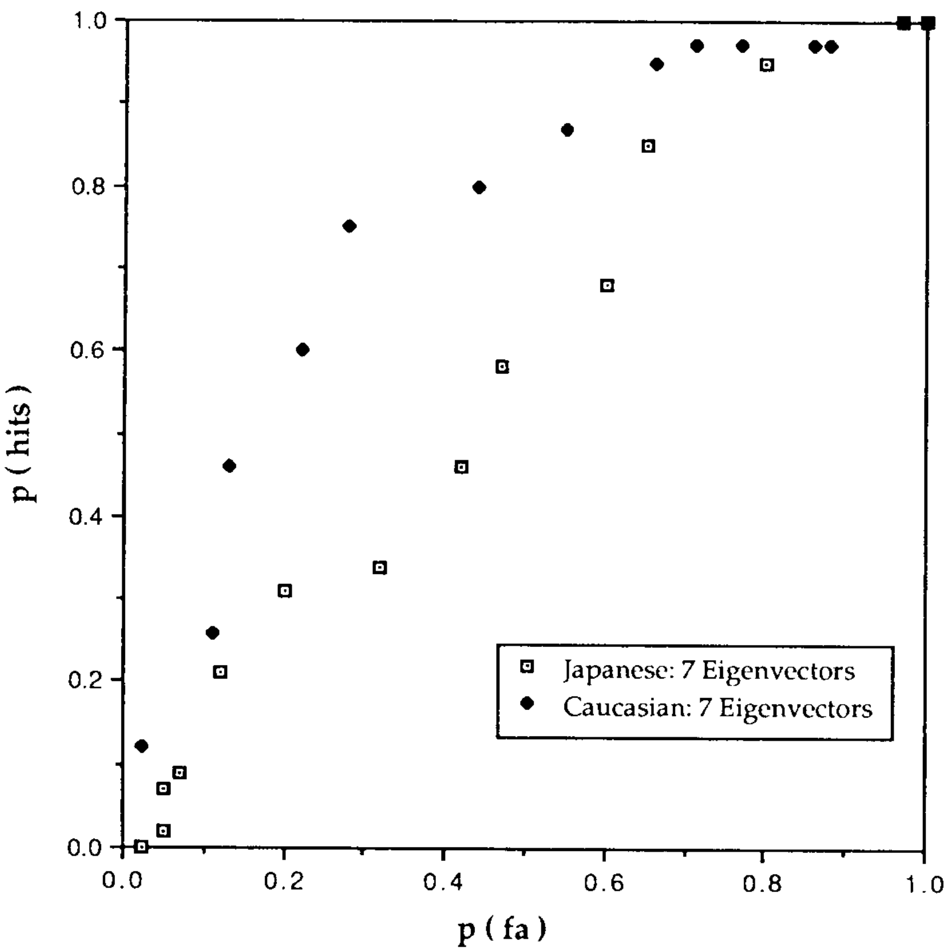

Figure 4. ROC curves for the long-term experience and short-term recognition Caucasian majority matrix.

race faces were used in these simulations and that race was probably the largest category difference in these simulations and is, therefore, likely to be represented in the first few eigenvectors.

We then simulated the short-term component by recalling faces combining the eigenvectors from the long-term majority matrix and the short-term half-Caucasian $(n = 40)$ and half-Japanese $(n = 40)$ matrix, weighting the long-term component at 0.75 and the short-term component at 0.25 [5]. We report a simulation only for the Caucasian majority matrix[6]. Here we see the classic cross-race effect, with the Japanese faces being more difficult to recognize (i.e. to separate OLD from NEW in the short-term recognition task) than the Caucasian faces. The ROC curves for this simulation are displayed in Figure 4.

The eigenvectors as features. Recalling faces from the autoassociative matrix is carried out by summing together a weighted combination of eigenvectors. That is, the faces are 'put together' by adding up the eigenvectors in differentially weighted combinations. As such, by most psychological definitions, the eigenvectors can be thought of as features of the faces. This interpretation of eigenvectors in associative matrices has been pointed out by Anderson et al. (1977). Also, in the context of low-dimensional representation of images, Sirovitch & Kirby (1987) suggest an

[5] These numbers are arbitrary and are simply an attempt to give more weight to the long-term experience than the short-term recognition task.

[6] This is because we do not yet have a sufficient number of Japanese faces available to complete the analysis for the Japanese faces.

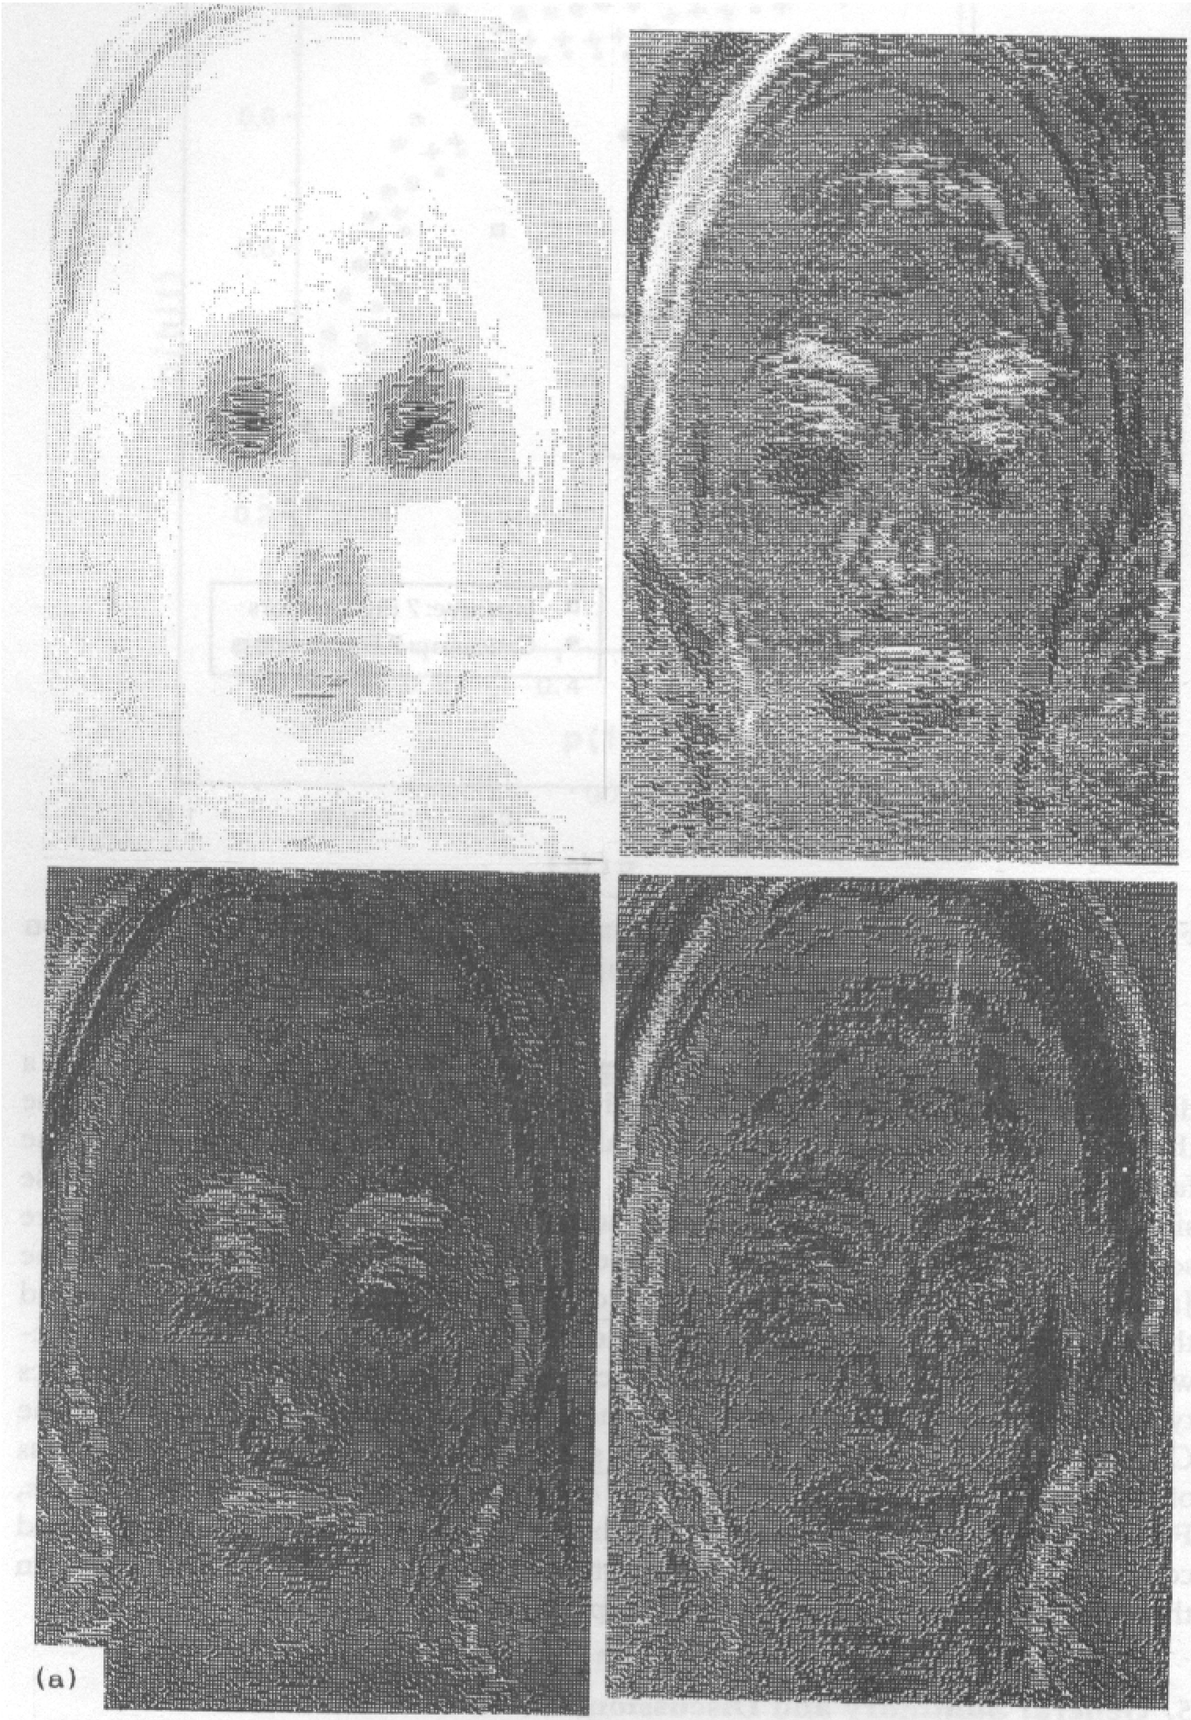

Figure 5. **(a)** The first four eigenvectors for the Caucasian majority simulation.

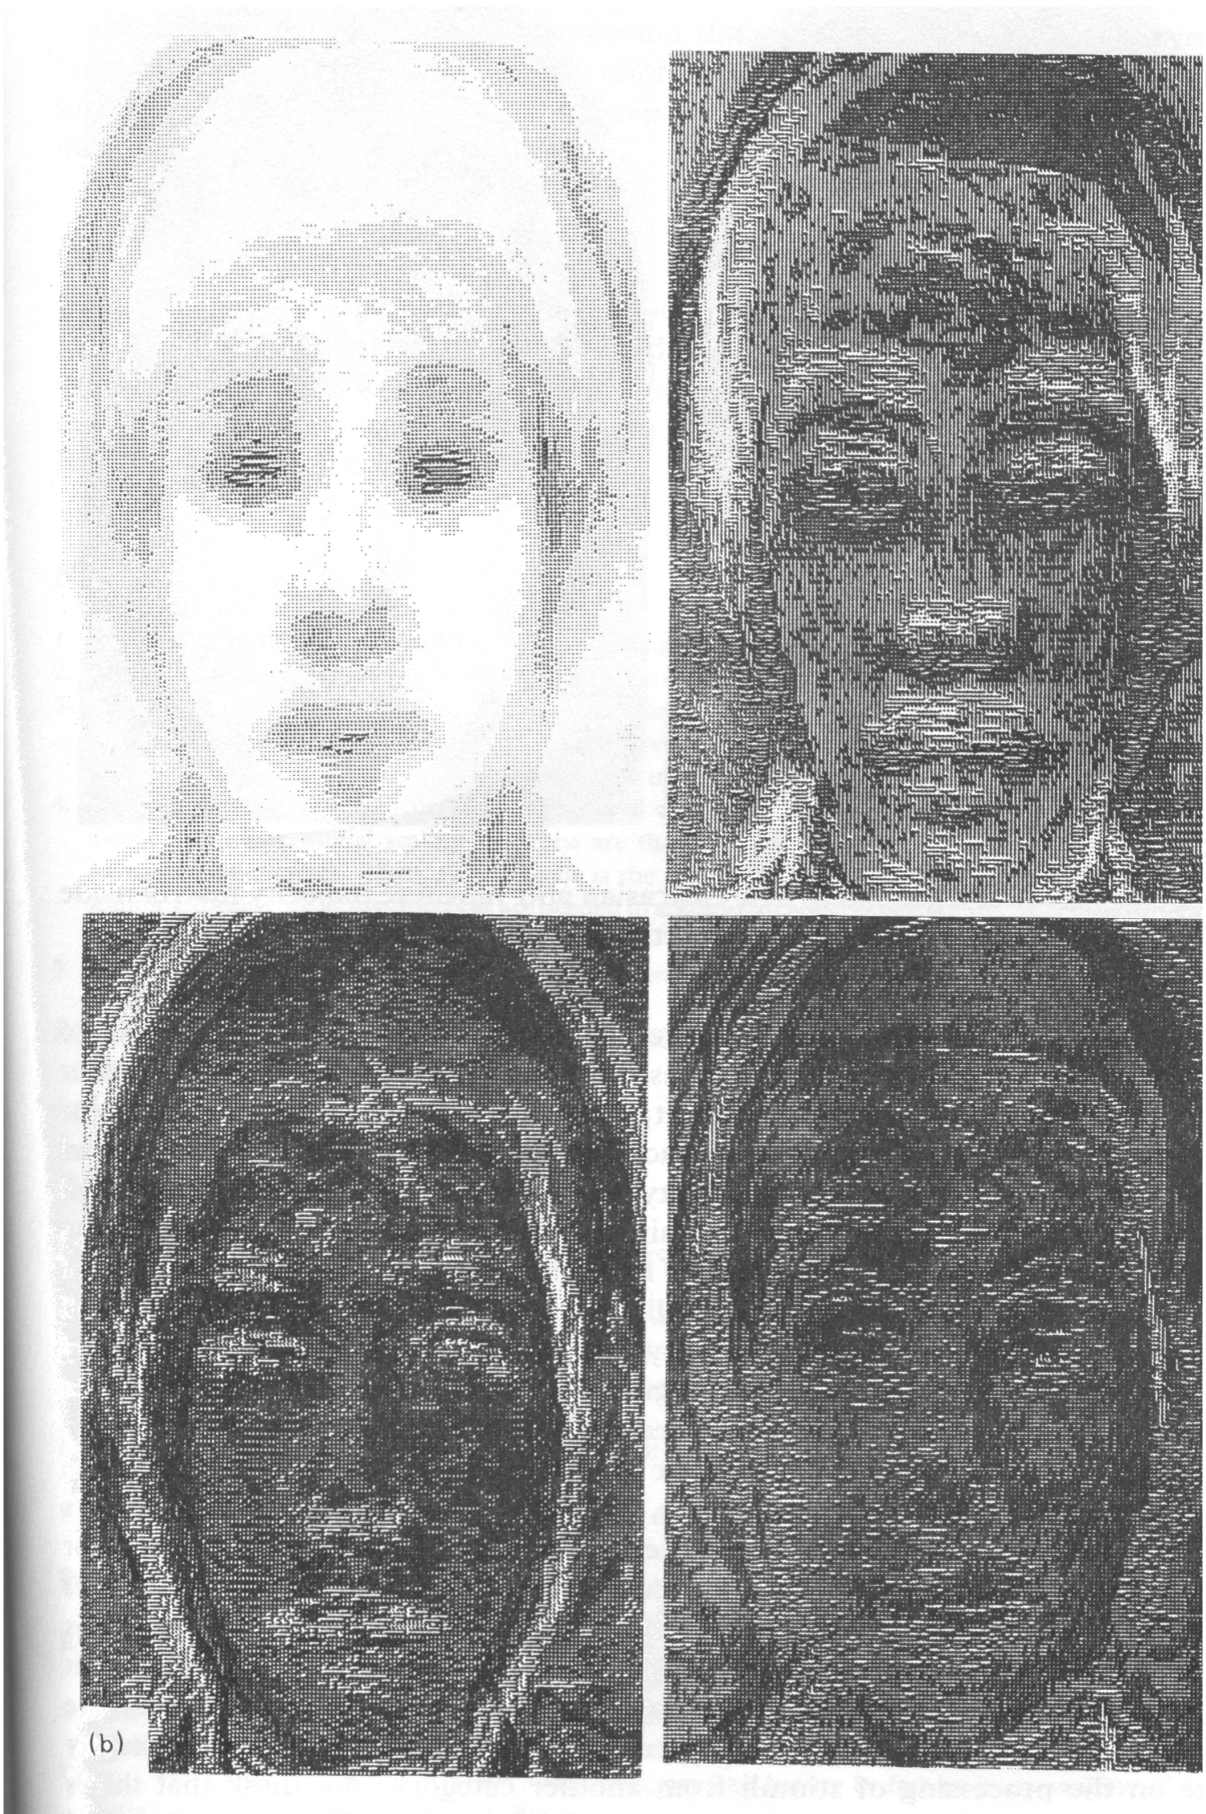

Figure 5. **(b)** (b) The first four eigenvectors for the Japanese majority simulation.

eigenvector-based description.

Applied to the current work, the eigenvectors of a matrix of face images are a different sort of feature than has generally been used in describing faces. For one thing, the eigenvectors represent global and not local features, since they span the face. Secondly, with the exception of the first eigenvector in each of these simulations, the eigenvectors are not readily interpretable in a traditional feature sense. The first four eigenvectors for the Caucasian majority simulation and the Japanese majority simulation are displayed in Figure 5 (a and b). It should be noted that the eigenvectors are face-like. Furthermore, the eigenvectors resemble somewhat the majority race of the matrix. The first eigenvector contains characteristics typical of the majority race (e.g. note the roundness of the eyes and face in the Caucasian majority eigenvectors, and the squareness of the face and distinctiveness of the nostrils for the Japanese majority eigenvectors). Finally, for completeness, Figure 6 shows the first eigenvector of each majority matrix made from a pixel-based code without spatial differentiation. The race differences are even more striking in these cases since shading information is preserved.

General Summary and Discussion

The purpose of these simulations was to model some common effects associated with processing other-race faces. We have tried to show that these effects can be modeled, in part, as a process of fine tuning to the information most useful for distinguishing faces within a homogenous set (i.e. a single race of faces). This tuning is suboptimal for processing other-race faces, however, and the system shows a number of shortcomings for the minority faces as compared to the majority faces. Our simulations produced three results. First, when the face history of a network was strongly biased toward a single race of faces, the model's ability to represent novel faces from this race exceeded its ability to represent faces from another race. Secondly, an autoassociative network trained on a majority race produced codings that were more similar to one another for faces of the minority race than for faces of the majority race. This simulates the well-known effect of faces of another race all appearing similar to one another. Finally, by combining a long-term face history experience matrix with a short-term recognition matrix, we simulated the other-race effect with majority faces being better recognized than minority faces.

While the system produced a number of effects that are qualitatively similar to those seen in the psychological literature, we caution that this approach is perhaps best thought of, not as a model of face recognition, but as an exploratory tool for quantifying and processing subtle perceptual information in complex images such as faces. It is also, not the only approach to simulating the other-race effect. Used in this context, it provides a method for examining other kinds of codings that might account for these effects in a similar fashion. Furthermore, its application might give insight into the constraints that extensive experience with a given stimulus category place on the processing of stimuli from another category. We think that this is especially important in cases where it is difficult to quantify the subtle visual information that separates the categories.

Finally, we think the model also has potential as a tool for simulating some other well-known effects in face memory, such as the relationship between typicality and recognition memory. Furthermore, it might be useful for giving insight into the perceptual components of the recognition difficulties encountered with inverted faces and with faces presented in the photographic negative. For these effects, it is instructive to pursue some simple perceptual explanations before looking at other more complicated explanations.

References

Abdi, H. (1988). A generalized approach for connectionist auto-associative memories: interpretation, implications and illustration for face processing. In J. Demongeot (Ed.) *Artificial Intelligence and Cognitive Sciences.* Manchester: Manchester University Press.

Anderson, J. A., Silverstein, J. W., Ritz, S. A. & Jones, R. S. (1977). Distinctive features, categorical perception and probability learning: some

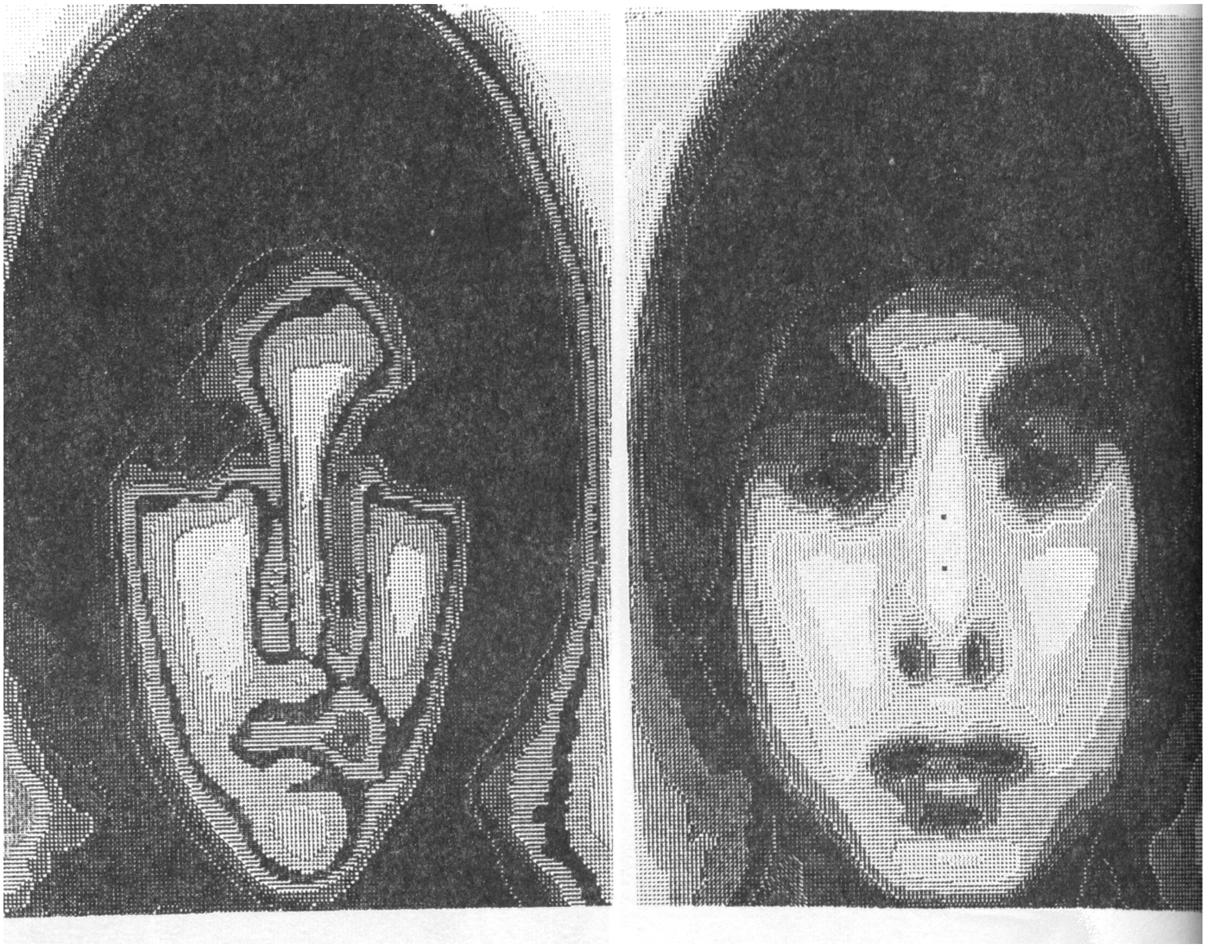

Figure 6. The first eigenvector of a Caucasian and Japanese majority matrix made from a pixel-based code without spatial differentiation.

applications of a neural model. *Psychological Review, 84*, 413–451.

Bothwell, R. K., Brigham, J. C. & Malpass, R. S. (1989). Cross-racial identification. *Personality and Social Psychology Bulletin, 15*, 19–25.

Brigham, J. C. (1986) The influence of race on face recognition. In H. D. Ellis, M. A. Jeeves, F. Newcombe & A. Young (Eds) *Aspects of Face Processing*. Dordrecht: Martinus Nijhoff.

Brigham, J. C. & Barkowitz, P. (1978). Do "they all look alike?" The effect of race, sex, experience, and attitudes on the ability to recognize faces. *Journal of Applied Social Psychology, 8*, 306–318.

Brigham,J. C., Malpass, A., Snyder, L. S. & Spaulding, K. (1982). The accuracy of eyewitness identification in a field setting. *Journal of Personality and Social Psychology, 42*, 673–681.

Chance, J. E., Turner, A. L. & Goldstein, A. G. (1982). Development of differential recognition of own-and other-race faces. *Journal of Psychology, 112*, 29–37.

Cross, J. G., Cross, J. & Daly, J. (1971). Sex, race, age and beauty as factors in recognition of faces. *Perception & Psychophysics, 10*, 393–396.

Deffenbacher, K. A. & Loftus, E. F. (1982). Do jurors share a common understanding concerning eyewitness behavior? *Law and Human Behavior, 6*, 15–30.

Feingold, C. A. (1914). The influence of environment on identification of persons and things. *Journal of Criminal Law and Police Science, 5*, 39–51.

Feinman, S. & Entwistle, D. R. (1976). Children's ability to recognize other children's faces. *Child Development, 47*, 506–510.

Goldstein, A. G. & Chance, J. E. (1985). Effects of training on Japanese face recognition: Reduction of the other-race effect. *Bulletin of the Psychonomic Society*, *23*, 211–214.

Hopfield, J. J. (1984). Neurons with graded responses have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, *81*, 3088–3092.

Kohonen, T. (1984). *Self Organization and Associative Memory*. Berlin: Springer Verlag.

Malpass, R. S. (1981) Training in face recognition. In G. Davies, H. Ellis & J. Shepherd (Eds) *Perceiving and Remembering Faces.* London: Academic Press, pp.271-285.

O'Toole, A. J. & Abdi, H. (1989). Connectionists approaches to visually based feature extraction. In G. Tiberghien (Ed.) *Advances in Cognitive Psychology, Vol. 2.* London: Wiley.

O'Toole, A. J., Millward, R. B. & Anderson, J. A. (1988). A physical system approach to recognition memory for spatially transformed faces. *Neural Networks*, *1*, 179–199.

Shapiro, P. N. & Penod, S. D. (1986). Meta-analysis of face identification studies. *Psychological Bulletin*, *100*, 139–156.

Sirovitch, L. &. Kirby, M. (1987). Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America*, *3*, 519–524.