The Karhunen-Loève Procedure for Gappy Data

R. Everson and L. Sirovich

The Rockefeller University

New York, NY 10021-6399

rme@camelot.rockefeller.edu *chico@camelot.rockefeller.edu*

November 1994

Abstract

This paper addresses the problem of using the Karhunen-Loève transform with partial data. Given a set of empirical eigenfunctions we show how to recover the modal coefficients for each gappy snapshot by a least-squares procedure. This method gives an unbiased estimate of the data that lay in the gaps and permits gaps to be filled in a reasonable manner. In addition, a scheme is advanced for finding empirical eigenfunctions from gappy data. It is shown numerically that this obtains spectra and eigenfunctions that are close to those obtained from unmarred data.

1 Introduction

A main purpose of this paper is to address the following question: How much image information is necessary for the restoration of a full image from a partial image, if it is known that the image belongs to a certain well-defined class of images? (Alternatively, how much degradation, by deletion of pixels, can such an image suffer and still be recovered?) Such questions are prompted by a number of applications in which image information is collected as an ensemble of *like* images and due to technical or natural circumstances some or all of the images are marred by gaps in the data. Many examples of this sort occur for data gathered from remote sensing satellites. As an illustration we mention the presence of cloud cover as a natural obstruction which leaves gaps in data records[1]. Although the language and illustrations presented here come from image analysis, the methodologies apply to the wider arena of databases having support in higher dimensions. While our deliberations may be relevant to image compression, this is not pursued here.

We first address the problem of recovering a full image from a marred image when the properties of an ensemble of like images are known. The methods rely on the Karhunen-Loève expansion for the ensemble and in the second part of the article we address the problem of finding the Karhunen-Loève expansion from a marred ensemble.

In order to deal with the issues involved we reconsider the *Rogues Gallery* problem which was formulated and solved in references [2] and [3]. Briefly stated this is the problem of analyzing an ensemble of images of human faces. A snapshot of a face will be denoted by $\phi = \phi(\mathbf{x})$ where ϕ represents deviation in gray level from the ensemble mean gray level at *pixel* location $\mathbf{x} = (x, y)$. If the faces are indexed by n, the ensemble is denoted by $\{\phi_n\}$, where $1 \leq n \leq N$ and N represents the number of faces in the ensemble. Normalization and other related details may be found in [2, 3]. It was shown there that there exists an optimal representation in the sense that the average error

$$\epsilon = \langle \| \phi - \sum_{n=1}^{M} a_n \psi_n(\mathbf{x}) \|^2 \rangle \tag{1}$$

is minimal for all M. Here, $\| \cdot \|^2$ denotes the usual L^2 norm and $\langle \cdot \rangle$ denotes the average over the ensemble. The minimal value of ϵ is obtained if the basis elements, ψ_n, satisfy the eigenfunction problem,

$$\int K(\mathbf{x},\mathbf{y})\psi_n(\mathbf{y})d\mathbf{y} \;=\; \lambda_n\psi_n(\mathbf{x}) \tag{2}$$

$$(\psi_n,\psi_m) \;=\; \int \psi_n(\mathbf{x})\psi_m(\mathbf{x})d\mathbf{x} \;=\; \delta_{nm} \quad , \tag{3}$$

where

$$K(\mathbf{x},\mathbf{y}) \;=\; \langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle \;=\; \frac{1}{N}\sum_{n=1}^{N}\phi(\mathbf{x})\phi(\mathbf{y}) \tag{4}$$

is the two point correlation function. This is the essence of the Karhunen-Loève (K-L) procedure or Principal Components Analysis (PCA), also known under a variety of other designations, and yields to standard numerical procedures. In what follows $\{\psi_n(\mathbf{x})\}$, which appear as eigenfunctions of the correlation operator, will be referred to as the empirical eigenfunctions. The K-L procedure, and variations of it, have been rediscovered a number of times. A very modern form of it goes back to Schmidt[4]. (For a description and extension see Sirovich and Everson[5].) Stewart[6] has recently reviewed the history of the method and a review of its use in turbulence theory appears in Berkooz *et al.*[7] Turk and Pentland[8] have subsequently made similar calculations to those presented in [2, 3] and use empirical eigenfunctions for face recognition. O'Toole and Abdi have used empirical eigenfunctions to investigate perception of race, sex and related studies[9, 10]. Kelly [11] has used a scheme similar to that advanced here for estimating modal coefficients of gappy data in an oceanographic context.

2 Marred Faces

With the use of the empirical eigenfunctions, ψ_n, which here we may call *eigenfaces*, only a relatively small number of parameters enter into the specification of a particular face. In quantitative terms it was found that on average fifty eigenfaces account for about 93% of the variance based on departures from the mean. [2] for a description of the ensemble and normalization. This should be compared with the $O(10^4)$ gray levels required to specify each snapshot.

This implies that for some face, $\phi(\mathbf{x})$, a suitable approximation can be obtained from a limited summation,

$$\phi(\mathbf{x}) \;\approx\; \sum_{n=1}^{N} a_n \psi_n(\mathbf{x}), \tag{5}$$

where the coefficients, a_n, are obtained from the usual inner product,

$$a_n \;=\; (\phi, \psi_n) \tag{6}$$

and N represents the number of basis functions needed to meet some specified error bound. Equation 5, looked at in another way, states that in the presence of perfect information (zero noise) we only need know the gray levels, $\phi(\mathbf{x})$, at N pixel locations.

To investigate this assertion we will consider *marred* faces and then investigate how well they can be reconstructed. We express a masked face by

$$\tilde{\phi}(\mathbf{x}) \;=\; m(\mathbf{x})\phi(\mathbf{x}) \tag{7}$$

where $m = 0$ on the mask and $m = 1$ elsewhere. The challenge is to write $\tilde{\phi}(\mathbf{x})$ in the form 5

$$\tilde{\phi}(\mathbf{x}) \;\approx\; m(\mathbf{x}) \sum_{n=1}^{N} \tilde{a}_n \psi_n(\mathbf{x}) \tag{8}$$

and from this determine a best set of coefficients $\tilde{a}_n$. Once this is done we can inquire as to how well ϕ is captured by $\sum_{n=1}^{N} \tilde{a}_n \psi_n$. Part of the problem involves the choice of N.

The inner product (6) can no longer be used to find the coefficients, because it requires information from the full range of $\mathbf{x}$, *i.e.* the ϕ_n are not necessarily orthogonal over the support of $\tilde{\phi}$, $s[\tilde{\phi}]$ However, we can then use a least-squares criterion to achieve a best fit of the form (5). That is, we minimize the error

$$E = \int_{s[\tilde{\phi}]} d\mathbf{x} \; \left(\tilde{\phi}(\mathbf{x}) - \sum_{n=1}^{N} \tilde{a}_n \psi_n \right)^2 \tag{9}$$

The minimization of E leads to

$$\left(\tilde{\phi} - \sum_{n=1}^{N} \tilde{a}_a \psi_n, \; \psi_k \right)_{s[\tilde{\phi}]} = 0 \tag{10}$$

which requires that the residual be orthogonal to ψ_k for $k = 1, \ldots, N$, where as indicated the inner product is over the support of $\tilde{\phi}$, $s[\tilde{\phi}]$. The Hermitian matrix

$$M_{kn} = (\psi_k, \psi_n)_{s[\tilde{\phi}]} \tag{11}$$

is non-negative and in principle $O(N)$.

If we write

$$f_k = (\phi, \psi_k)_{s[\tilde{\phi}]} \tag{12}$$

then in vector notation, we seek the unknown coefficients $\tilde{a}_k$ from

$$\mathbf{M}\tilde{\mathbf{a}} = \mathbf{f} \tag{13}$$

In the event that $s[\tilde{\phi}]$ is sufficiently dense in the space then $\mathbf{M} \approx \mathbf{I}$, which among other properties, says that the eigenvalues of $\mathbf{M}$ are close to unity, and $\tilde{a}_k \approx (\phi, \psi_k)_{s[\tilde{\phi}]}$. In the present instance, if we denote the eigenvalues by μ_n and the corresponding orthonormal eigenvectors by $\mathbf{v}_n$ the solution to (13) is then given by

$$\tilde{\mathbf{a}} = \sum_{k=1}^{N} \frac{1}{\mu_n} (\mathbf{v}_n, \mathbf{f}) \mathbf{v}_n \tag{14}$$

Thus, on intuitive grounds, the construction becomes questionable if the μ_k depart significantly from unity; this is made exlicit in the Appendix.

To illustrate the nature of this construction consider the mask shown in Figure 1a. This is a relatively extreme mask which obscures ninety percent of the pixels in a randomly chosen way. This was used to mask a face, not belonging to the original ensemble used to determine the eigenfunctions. The result of applying the above procedure, finding the $\tilde{a}$ from (13) and using $N = 50$ eigenfunctions, is shown in Figure 1b. The original unmasked face is shown in Figure 1c and the projection of the original face onto 50 eigenfunctions is shown in Figure 1d. Though the procedure does not recover the original face exactly, the construction is visually very close to the projection onto 50 eigenfunctions which utilizes the entire area and is the best that may be achieved with 50 functions.

We underline the fact that the masked face did not enter into the determination of the eigenfaces. When the face to be reconstructed is a member of the ensemble used to construct the eigenfunctions both the reconstructed face and the projected face are closer to the original. Figure 2 show the result of a reconstruction for a face that was a member of the original ensemble.

In carrying out this construction we have taken $N = 50$ in (10). This, as Figure 3 shows, is an optimal choice for the number of fitting functions when the fraction of unobscured pixels $p = 0.1$. Figure 3 shows the mean squared error, $\int |\phi - \sum_{n=1}^{N} \tilde{a}\phi|^2 d\mathbf{x}$, averaged over 48 faces not part of the ensemble. The bottom curve corresponds to $p = 1$. Here the entire face is unmasked, the coefficients are determined by the simple inner product and the error is the best that may be attained for a particular N, though it is necessarily larger than $\sum_{n=N+1} \lambda_n$. When $p \geq 0.2$ the least-squares procedure performs well for all N.

In Figure 4 we show the mean squared error versus N averaged over the 238 faces comprising the original ensemble. The errors are significantly smaller here than for the previous case, reflecting the fact that the empirical eigenfunctions are optimally suited to this particular ensemble. Again, lower curve corresponds to $p = 1$, for which the entire face is unmasked, the coefficients are determined by the simple inner product and the error is the best that may be attained for a particular N. In fact, this best error is given by $\sum_{n=N+1} \lambda_n$. Clearly, when $p \geq 0.05$ the least-squares procedure performs well.

The basic reason for the ability of the procedure just presented to recover the marred regions depends, in part, on the fact that only a limited number of fitting functions are needed in order to well approximate a full face, suitably normalized and satisfying other reasonable requirements. This number, which might appropriately be called, the dimension of *face space*, is roughly 50. This estimate is greatly dependent on the use of the empirical eigenfunctions, which span face space in an optimal manner. Other basis functions can require many more fitting functions.

If we take this as a nominal value, then only 50 coefficients a_n need be determined in (10). Since $O(500)$ pixels are not masked by the 10% mask it is clear why the least-squares fit leads to a successful answer. It is at least intuitively clear that a sufficiently masked image cannot correctly furnish the smaller scales. Since the eigenfunctions resolve successively smaller scales with increasing index, N, using *too many* eigenfunctions results in a deterioration of the fit. This behavior results from a trade-off between the possibility of a better fit using more eigenfunctions and the fact that $\mathbf{M}$ increasingly departs from the identity as the amount of information (number of pixels) available to determine each coefficient decreases. When the unmasked area is sufficiently large ($p > 0.05$ when the face belongs to the original ensemble and $p > 0.2$ when it does not) using more eigenfunctions always results in a better fit. A more detailed analysis and estimates for the optimum N given p are presented in the Appendix.

One further aspect of the analysis merits comment. It is important that the image which lies beneath the mask be a member of the class from which the empirical eigenfunctions were obtained; in this instance eigenfunctions were obtained from an ensemble of shaven, Caucasian males, though it is known that they are suitable for female faces. In Figure 5 we show the result of reconstructing a monkey face from under 50% mask. (The monkey face was scaled and normalized in exactly the same manner as the human faces.) The unobscured face is shown at the left; the reconstruction, using $n = 100$ eigenfunctions, from a $p = 0.5$ random mask in the middle, and the reconstruction using 220 eigenfunctions and the entire area at the right. The mean squared errors are approximately an order of magnitude larger for the monkey face than out-of-ensemble human faces. It is clear that the human eigenfunctions are unsuited to the monkey face and such wide departures from the class as bearded faces. The reconstruction can only select human components of the simian face. All of this implies that some law governs the organization of a human face. While the law remains unknown, our results imply that no more than $O(50)$ dimensions are needed for a reasonable characterization of *face space*.

3 Marred Eigenfunctions

Next we explore the determination of the eigenfunction set $\{\psi_n(\mathbf{x})\}$ if only marred data are available. For this purpose we consider an ensemble of masks, $\{m_n(\mathbf{x})\}$, with the value of each $m_n(\mathbf{x})$ either zero or unity depending on whether the pixel $\mathbf{x}$ is masked or not masked, respectively. The masks are randomly generated and we will characterize a mask by p, the fraction of unmasked pixels.

We denote the ensemble of marred faces by $\{\tilde{\phi}(\mathbf{x})\}$. Each marred face is of the form

$$\tilde{\phi}(\mathbf{x}) \;=\; m_n(\mathbf{x})\phi(\mathbf{x}) \tag{15}$$

where $\phi(\mathbf{x})$ is chosen from the original ensemble. Each face may occur more than once. If P denotes the total number of pixels then there exist 2^P possible masks, and the masks may be conveniently regarded as all different. Thus the ensemble $\{\tilde{\phi}(\mathbf{x})\}$ may be regarded as a significantly larger ensemble than $\{\phi(\mathbf{x})\}$.

The object of this section is to present an algorithm for the construction of the eigenfunctions and to show the results of this procedure. We will be content here to demonstrate existence, convergence and related questions by means of numerical investigations.

To start what will become an iterative procedure we define the average value at pixel location $\mathbf{x}$ by

$$\langle\tilde{\phi}(\mathbf{x})\rangle \;=\; \frac{1}{M(\mathbf{x})}\sum_{n\epsilon S[\mathbf{x}]}\tilde{\phi}_n(\mathbf{x}), \tag{16}$$

where $S[\mathbf{x}]$ is the set of indices at which $m_n(\mathbf{x})$ is unity and $M(\mathbf{x})$ is the number of indices in this set for pixel location $\mathbf{x}$. This average is simply the average over all the available information at the location $\mathbf{x}$.

As a first step of an iterative procedure each $\tilde{\phi}_n(\mathbf{x})$ is *repaired* by filling in missing pixels by the average values at those locations. We denote this repaired ensemble by $\{\tilde{\phi}_n^{(0)}(\mathbf{x})\}$. Since this ensemble is defined everywhere we can employ the Karhunen-Loève procedure to generate $\{\psi_n^{(0)}(\mathbf{x})\}$, a complete orthonormal system. Next we obtain $\{\tilde{\phi}_n^{(1)}(\mathbf{x})\}$ by fitting each $\tilde{\phi}_n(\mathbf{x})$ of the original ensemble by a superposition of R eigenfunctions $\{\psi_n^{(0)}(\mathbf{x})\}$ as follows: Set

$$\hat{\phi}^{(1)} = \sum_{n=1}^{R} a_n^{(1)}\psi_n^{(0)}(\mathbf{x}) \tag{17}$$

and determine the set $\{a_n^{(1)}\}$ by minimizing the criterion function over the pixels where data is available; that is, minimize

$$\tilde{E}_n = \int (\tilde{\phi}_n - \hat{\phi}_n^{(1)})^2 m_n(\mathbf{x}) d\mathbf{x} \tag{18}$$

The repaired snapshot, $\tilde{\phi}_n^{(1)}$ is now obtained by filling in the masked pixels with the $\hat{\phi}_n^{(1)}$:

$$\tilde{\phi}_n^{(1)}(\mathbf{x}) = \begin{cases} \tilde{\phi}_n(\mathbf{x}) & \text{if } m_n(\mathbf{x}) = 1 \\ \hat{\phi}_n^{(1)}(\mathbf{x}) & \text{if } m_n(\mathbf{x}) = 0 \end{cases} \tag{19}$$

This procedure is carried out for each $\tilde{\phi}_n$ and each $m_n(\mathbf{x})$.

The set $\{\tilde{\phi}_n^{(1)}\}$ is now defined everywhere and hence through K-L generates $\{\psi_n^{(1)}(\mathbf{x})\}$, an orthonormal complete system. The iteration is now clear and we next discuss the results.

4 Results

The iteration scheme is demonstrated on an ensemble of 286 faces masked so that 40% of the face is obscured. A typical example, denoted $\tilde{\phi}_{53}(\mathbf{x})$ is shown in Figure 6. Each mask consists of a union of squares with randomly distributed centers. The width of each square was drawn from a Poisson distribution with mean width 3 pixels.

Also shown in Figure 6 are the intermediate snapshots $\tilde{\phi}_{53}^{(1)}$, $\tilde{\phi}_{53}^{(2)}$, $\tilde{\phi}_{53}^{(5)}$, $\tilde{\phi}_{53}^{(10)}$ and $\tilde{\phi}_{53}^{(20)}$ as the iteration proceeds. At each stage the faces were repaired with $R = 30$ eigenfunctions $\psi_n^{(k)}$ (cf. equation 17). This R was judged to be large enough to capture the essential features of "face space", but not large enough to prohibit accurate determination of the $\tilde{a}_n^{(i)}$.

The convergence of the eigenvalue spectrum is illustrated in Figure 7, which shows the principal 60 eigenvalues, λ_n, after 1, 2, 10 and 20 iterations, together with the spectrum derived from the unmarred ensemble. The dominant eigenvalues display a clear convergence towards those of the unmarred spectrum. The "step" at index 30 is a consequence of repairing the faces with $R = 30$ eigenfunctions. In fact, since only 30 eigenfunctions are used one cannot hope to achieve better eigenfunctions and eigenvalues than those produced by repairing (using the scheme outlined above) the marred data with 30 eigenfunctions derived from perfect data. We denote these eigenfunctions and eigenvalues by χ_n and ρ_n. Figure 8 compares the spectra from the unmarred faces (λ_n), from the marred faces repaired with 30 unmarred eigenfunctions (ρ_n) and from iteration 20 of the scheme ($\lambda_n^{(20)}$). The iterative scheme, which lacks information about the perfect eigenfunctions, well approximates the initial portion of the other two spectra.

Although the eigenvalues are in good agreement, it remains to be checked that the eigenfunctions from the iterative scheme approximate the eigenfunctions from the unmarred data. Of relevance here is not the convergence of individual eigenfunctions, but of the spaces spanned by groups of eigenfunctions. Assessing the convergence of the eigenfunctions therefore requires a method of comparing subspaces. Let $\mathbf{E}$ and $\mathbf{F}$ be the projectors defining a pair of subspaces, each of dimension d. Then the trace of the Hermitian matrix, $\mathrm{Tr}\mathbf{EFE}$ measures the commonality of the subspaces. If the subspaces are identical $\mathrm{Tr}\mathbf{EFE} = d$; if they are disjoint $\mathrm{Tr}\mathbf{EFE} = 0$. We denote by $C_d(e_n, f_n)$ the trace of $\mathbf{EFE}$, where $\mathbf{E}$ and $\mathbf{F}$ are projectors for the subspaces spanned by the collections of vectors e_n and f_n. The commonality between the subspaces spanned by the first 30 unmarred eigenfunctions and the eigenfunctions derived from the faces repaired with unmarred eigenfunctions is given by $C_{30}(\psi_n, \chi_n) = 29.4$.

Figure 9 shows $C_d(\psi, \tilde{\psi}^{(k)})$ versus d for iterations $k = 1$ & 20. The rate of convergence is shown in Figure 10 in which $C_{30}(\psi, \psi^{(1)})$ and $C_{30}(\chi, \psi^{(1)})$ are plotted against iteration number. Even after a single iteration there is considerable overlap between the first 7 eigenfunctions. The rate of convergence is initially rapid, slowing as the limit is approached. After 20 iterations there is a very good agreement between the eigenfunctions from the iteration scheme and the unmarred eigenfunctions; visually, they are indistinguishable. Pursuing the iteration further improves the match, but the rate of convergence is slow; after 80 iterations $C_{30}(\psi, \psi^{(1)}) = 26.51$, which is to be compared with 26.26 after 20 iterations. It is possible that Newton's method or other schemes in which more eigenfunctions are introduced as the iteration proceeds would enhance the rate. We remark that $\psi_n^{(k)}$ approach the χ_n more closely and more rapidly.

5 Summary

We have addressed the problem of using the Karhunen-Loève transform with partial data. Given a set of eigenfunctions we have shown how to recover the modal coefficients for each gappy snapshot by a least-squares procedure. This method gives an unbiased estimate of the data that lay in the gaps and permits that gaps to be filled in a reasonable manner. In addition we have advanced a scheme for finding empirical eigenfunctions from the gappy data and have shown numerically that it yields a spectrum and eigenfunctions that are close to those obtained from unmarred data.

Acknowledgements

This work was performed with support from NASA (NAG 5-2336), the Office of Naval Research and the Navy Research Laboratory at Stennis (N00014-93-1-G901). We are grateful to N. Hochman for furnishing the monkey image used in our investigation.

6 Appendix

This appendix gives a more thorough discussion of the errors inherent in fitting a face with N empirical eigenfunctions when a fraction p of the pixels are unobscured. We can express an unmasked face by a superposition of the empirical eigenfunctions:

$$\phi(\mathbf{x}) = \sum_n a_n \psi_n(\mathbf{x}) \quad . \tag{20}$$

If the number of faces in the ensemble equals the number of pixels describing a face then the $\psi_n(\mathbf{x})$ are complete (even if this is not true a complete set can always be determined). Since the purpose of this appendix is to explore the interplay of the fraction, p, of unmasked pixels with a truncation number, N, we do not dwell further on this point.

For the truncation N we write

$$\phi(\mathbf{x}) \;=\; \sum_{n=1}^{N} a_n \psi_n(\mathbf{x}) + r(\mathbf{x}) \tag{21}$$

where the residual r depends on N:

$$r(\mathbf{x}) \;=\; \sum_{n>N} a_n \psi(\mathbf{x}) \quad . \tag{22}$$

Note that when the face being fitted belongs to the ensemble from which the eigenfunctions were determined the mean size of the residual is given by $<\| r \|^2>= \sum_{n>N} \lambda_n$. Thus if a masked face is written as $\tilde{\phi} = m(\mathbf{x})\phi(\mathbf{x})$, then

$$\tilde{\phi} \;=\; \sum_{n=1}^{N} a_n m(\mathbf{x}) \psi_n(\mathbf{x}) + \tilde{r}(\mathbf{x}) \tag{23}$$

with

$$\tilde{r}(x) \;=\; \sum_{n>N} a_n m(\mathbf{x}) \psi(\mathbf{x}). \tag{24}$$

From this it follows that $\mathbf{f}$ in (13) can be expressed as

$$\mathbf{f} \;=\; \mathbf{M}\mathbf{a} + \mathbf{e}, \tag{25}$$

where $\mathbf{a}^{\dagger} = (a_1, \ldots, a_N)$,

$$e_k = \sum_{n>N} a_n (\psi_k, \psi_n)_{s[\tilde{\phi}]} \quad , \tag{26}$$

and $\mathbf{M}$ is given by (11). Thus from (13)

$$\tilde{\mathbf{a}} \;=\; \mathbf{a} + \mathbf{M}^{-1}\mathbf{e} \tag{27}$$

and the accuracy of the approximation depends on $\| \mathbf{M}^{-1} \|$, which in turn is measured by the smallest eigenvalue of $\mathbf{M}$.

The form of $\mathbf{M}$ is given by (11). As before, we denote the fraction of unobscured pixels by p. The total number of pixels is denoted by T. Since $(\psi_n, \psi_n) = 1$, we estimate $\psi = O(1/\sqrt{T})$ at a pixel location. The entries of $\mathbf{M}$ are determined by sums of pT terms. In the application at hand $T = O(10^4)$ and p is not smaller than $O(10^{-2})$. It is therefore reasonable to expect that the central limit theorem will apply to these sums.

From these preliminary considerations we can express the symmetric matrix $\mathbf{M}$ as

$$\mathbf{M} \ = \ p\mathbf{I} \ + \ \mathbf{S} \tag{28}$$

where $\mathbf{S}$ is symmetric and has entries which are Gaussian distributed and of mean zero. If s denotes an entry of $\mathbf{S}$, then

$$P(s) = \frac{1}{\sigma\sqrt{2\pi pT}} \exp \frac{-s^2}{\sigma^2 pT} \tag{29}$$

The variance, σ^2 is somewhat problematic. For relatively large index $\psi_n(\mathbf{x})$ is locally sinusoidal. If we adopt this as a hypothesis it then follows that $\sigma^2 = O(1/T^2)$, and (29) takes the form,

$$P(s) = \sqrt{\frac{T}{2\pi p}} \exp \frac{-s^2 T}{p} \tag{30}$$

This implies that the off-diagonal terms of $\mathbf{M}$ are $O(\sqrt{p/T})$, and hence that the ratio of off-diagonal to on-diagonal terms is $O(1/\sqrt{pT})$. A series of numerical experiments confirms the Gaussian nature of the entries of $\mathbf{S}$ and also that the estimate for the ratio is sound (see Figure 11 and Table 1).

According to Wigner's semi-circle theorem[12, 13], a symmetric random matrix of order N, whose entries have zero mean and variance σ^2 has the eigenvalue density

$$\rho(\lambda) = \frac{1}{2\pi N\sigma^2}\sqrt{4N\lambda^2 - \mu^2}. \tag{31}$$

The theorem applies to $\mathbf{S}$ and since eigenvalues of $\mathbf{M}$ are p plus the eigenvalues of $\mathbf{S}$ it follows that the eigenvalues of $\mathbf{M}$ are such that

$$p - 2\sqrt{\frac{pN}{T}} \ < \ \mu \ < \ p + 2\sqrt{\frac{pN}{T}}. \tag{32}$$

If we take the vanishing of the smallest eigenvalue of $\mathbf{M}$ as a criterion for the breakdown of the scheme, then this occurs for

$$N \ = \ \frac{pT}{4} \tag{33}$$

In figures 3 and 4 we display the errors incurred for various p for the cases of faces not belonging to the ensemble and belonging to the ensemble which determined the eigenfunctions ψ_n. Given the rough nature of the estimate the criterion (33) is not bad. The two cases differ as a result of the fact that $<\| r \|^2>$ is significantly smaller when the faces belong to the ensemble which determines the ψ_n.

References

[1] R.M. Everson, P. Cornillion, L. Sirovich, and A. Webber. An empirical eigenfunction analysis of sea surface temperatures in the north atlantic. *Journal of Physical Oceanography*, 1995.

[2] L. Sirovich and M. Kirby. Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America*, 4A(3):519–524, 1987.

[3] M. Kirby and L. Sirovich. Application of the karhunen-loÈve procedure for the characterization of human faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1), 1990.

[4] E. Schmidt. Zur theorie der linearen und nichtlinearen integralgleichungen. i teil: Entwicklung willkÜlicher funktion nach systemen vorgeschriebener. *Mathenatische Annalen*, 63:433–476, 1907.

[5] L. Sirovich and R.M. Everson. Analysis and management of large scientific databases. *International Journal of Supercomputing Applications*, 6(1):50–68, 1992.

[6] G.W. Stewart. On the early history of the singular value decomposition. *SIAM Review*, 35(4):551–566, 1993.

[7] G. Berkooz, P. Holmes, and J.L. Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Reviews Fluid Mechanics*, 25, 1993.

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

[9] A.J. O'Toole, H. Abdi, K.A. Deffenbacher, and J.C. Bartlett. Classifying faces by race and sex using an autoassiciative memory trained for recognition. In K.J. Hammond and D. Gentner, editors, *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 847–851. 1991.

[10] A.J. O'Toole, K.A. Deffenbacher, H. Abdi, and J.C. Bartlett. Simulating the "other-race effect" as a problem in perceptual learning. *Connection Science*, 3(2):163–178, 1991.

[11] K.A. Kelly. The influence of winds and topography on sea surface temperature patterns over the northern california slope. *Journal of Geophysical Research*, 90(C6):11783–11798, 1985.

[12] E.P. Wigner. Random matrices in physics. *SIAM Review*, 9(1), 1967.

[13] M. Carmeli. *Statistical Theory and Random Matrices.* Marcel Dekker, New York, 1983.

Figure Captions

Figure 1 Reconstruction of a face, not in the original ensemble, from a 10% mask. The reconstructed face (b) was determined using 50 empirical eigenfunctions and only the white pixels shown in (a). The original face is shown in (c) and a projection (using all the pixels) of the face onto 50 empirical eigenfunctions is shown in (d).

Figure 2 Reconstruction of a face from a 10% mask. The reconstructed face (b) was determined using 50 empirical eigenfunctions and only the white pixels shown in (a). The original face, which was a member of the ensemble is shown in (c) and a projection (using all the pixels) of the face onto 50 empirical eigenfunctions is shown in (d).

Figure 3 Mean squared error versus the number of fitting eigenfunctions for snapshots that were not part of the original ensemble. Different curves show the error for different unmasked areas, p. When $p = 1$ the entire picture is unmasked and the mean squared error is the best that may be attained for a given number of fitting eigenfunctions.

Figure 4 Mean squared error versus the number of fitting eigenfunctions for snapshots that were members of the original ensemble. Different curves show the error for different unmasked areas, p. When $p = 1$ the entire picture is unmasked and the mean squared error is the best that may be attained for a given number of fitting eigenfunctions.

Figure 5 Reconstruction of a monkey face using eigenfunctions derived from human faces. Left: unobscured face. Middle: reconstruction from a $p = 0.5$ random mask, using 100 eigenfunctions. Right: reconstruction using 220 eigenfunctions and the entire area.

Figure 6 A 40% masked face $\tilde{\phi}_{53}$and the intermediate snapshots $\tilde{\phi}_{53}^{(1)}$, $\tilde{\phi}_{53}^{(2)}$, $\tilde{\phi}_{53}^{(5)}$,$\tilde{\phi}_{53}^{(10)}$ and $\tilde{\phi}_{53}^{(20)}$ as the iteration scheme proceeds. At each stage the masked regions have been repaired with $R = 30$ eigenfunctions derived from the snapshots from the previous iteration.

Figure 7 Eigenvalue spectrum, λ_n, after 1, 2 and 20 iterations. The solid line shows the spectrum derived from the unmarred snapshots.

Figure 8 Eigenvalue spectrum, $\lambda_n^{(20)}$, after 20 iterations compared with the unmarred spectrum, λ_n, and the spectrum, ρ_n, derived from the snapshots repaired with eigenfunctions from the unmarred ensemble.

Figure 9 Commonality of subspaces spanned by the unmarred eigenfunctions, ψ_n, and the repaired eigenfunctions, ψ_n^k, after $k = 1$ and $k = 20$ iterations.

Figure 10 Convergence with iteration of the subspaces spanned by the repaired eigenfunctions and the unmarred eigenfunctions (squares) and the subspaces spanned by the repaired eigenfunctions and the eigenfunctions derived from marred data repaired with perfect eigenfunctions (triangles).

Figure 11 Empirical probability densities, $P(s)$, of the off-diagonal elements of the matrix $\mathbf{M}$, illustrating the Gaussian nature of the distribution. The matrices were generated at the following parameters: (a) $p = 0.7$, $N = 100$; (b) $p = 0.7$, $N = 50$; (c) $p = 0.1$, $N = 100$.

Table 1 Ratio of the mean size of the diagonal elements to the off-diagonal elements of the matrix $\mathbf{M}$, divided by $\sqrt{pT}$, for various values of p and N. Except when $p \approx 1$, the estimate that the ratio is $O(\sqrt{pT})$ is verified.

N	p	Normalized ratio
10	0.06	1.47
100	0.06	1.14
150	0.06	1.13
200	0.06	1.12
10	0.25	1.18
100	0.25	1.31
150	0.25	1.27
200	0.25	1.29
10	0.5	2.72
100	0.5	1.66
150	0.5	1.55
200	0.5	1.54
10	0.75	3.01
100	0.75	2.11
150	0.75	2.20
200	0.75	2.16
10	0.95	5.28
100	0.95	4.98
150	0.95	4.59
200	0.95	5.05

Table 1: Ratio of the mean size of the diagonal elements to the off-diagonal elements of the matrix $\mathbf{M}$, divided by $\sqrt{pT}$, for various values of p and N. Except when $p \approx 1$, the estimate that the ratio is $O(\sqrt{pT})$ is verified.

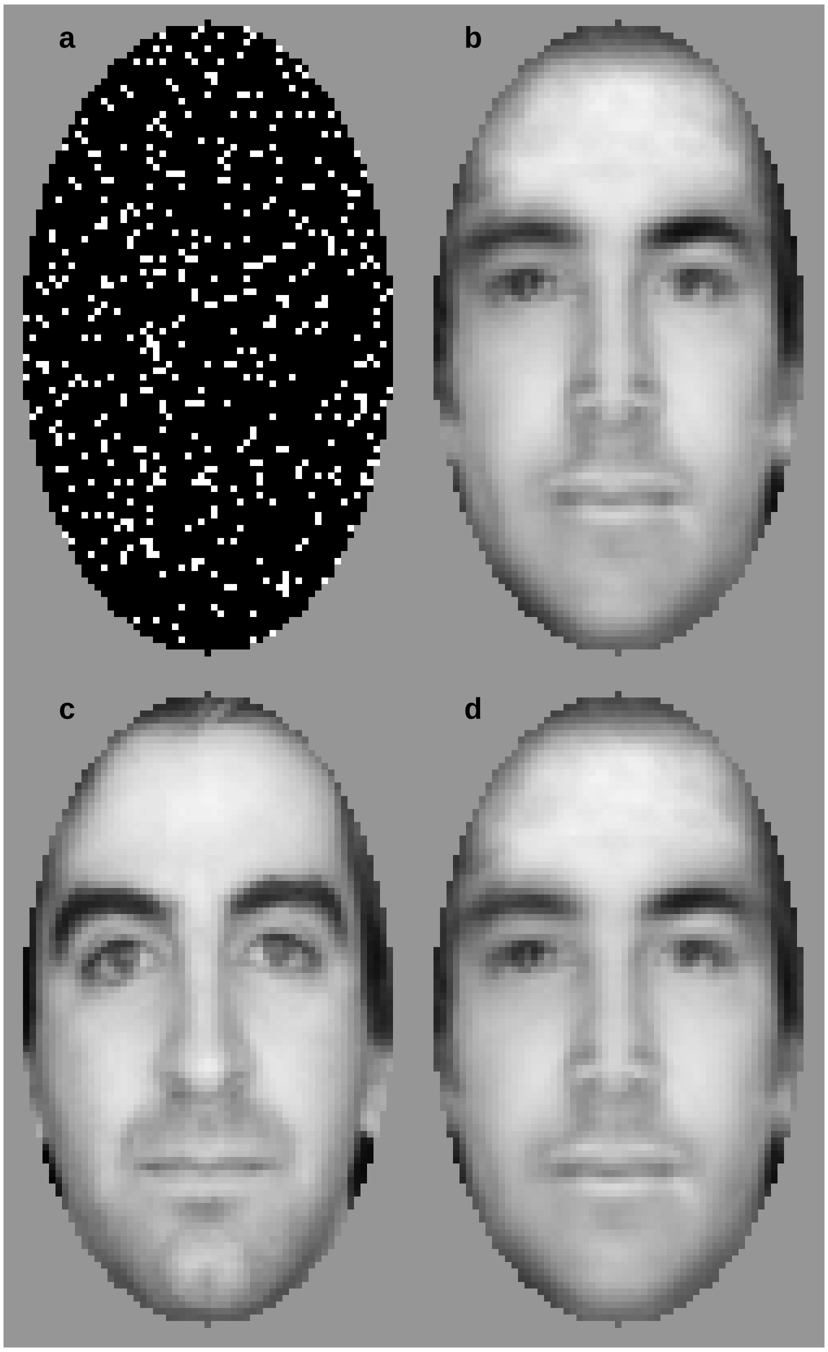

Figure 1: Reconstruction of a face, not in the original ensemble, from a 10% mask. The reconstructed face (b) was determined using 50 empirical eigenfunctions and only the white pixels shown in (a). The original face is shown in (c) and a projection (using all the pixels) of the face onto 50 empirical eigenfunctions is shown in (d).

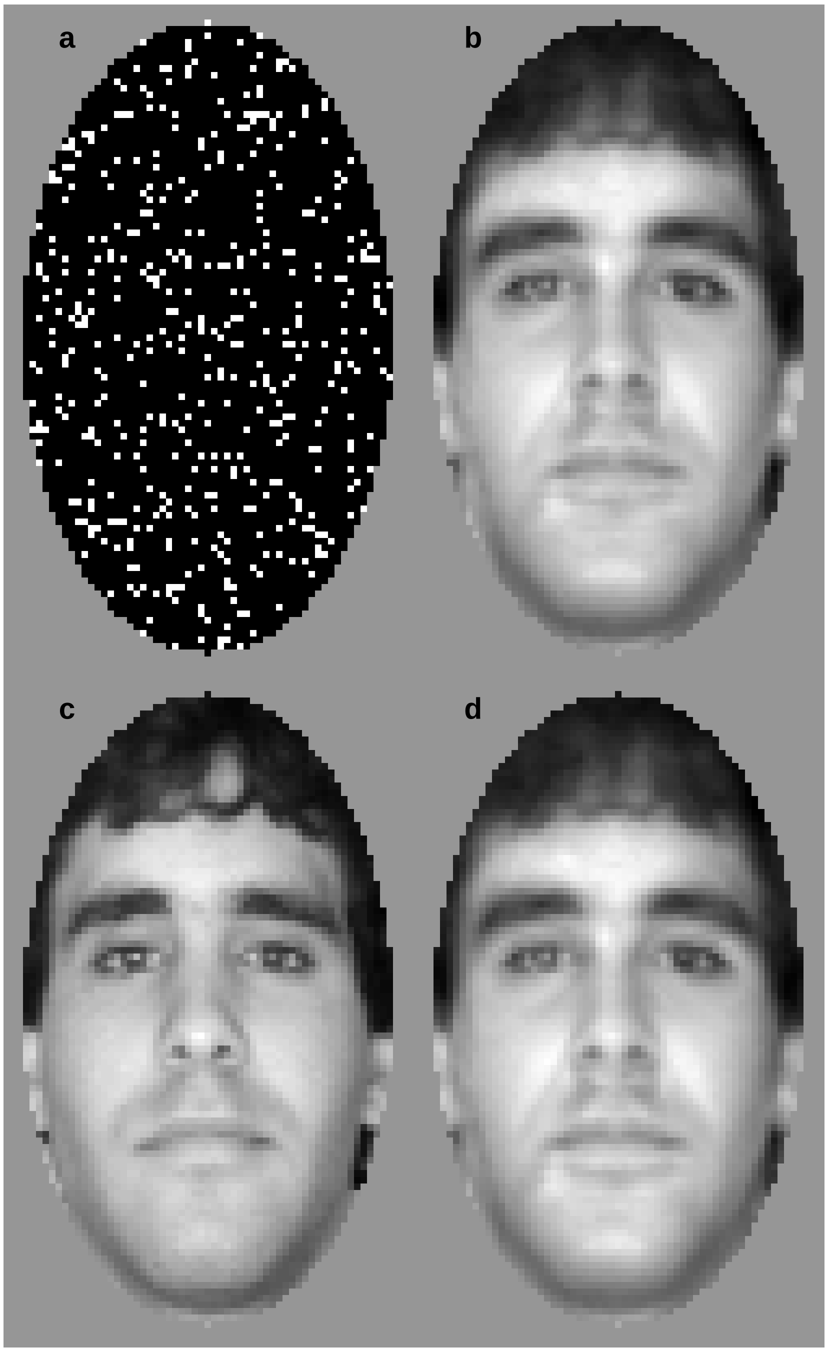

Figure 2: Reconstruction of a face from a 10% mask. The reconstructed face (b) was determined using 50 empirical eigenfunctions and only the white pixels shown in (a). The original face, which was a member of the ensemble is shown in (c) and a projection (using all the pixels) of the face onto 50 empirical eigenfunctions is shown in (d).

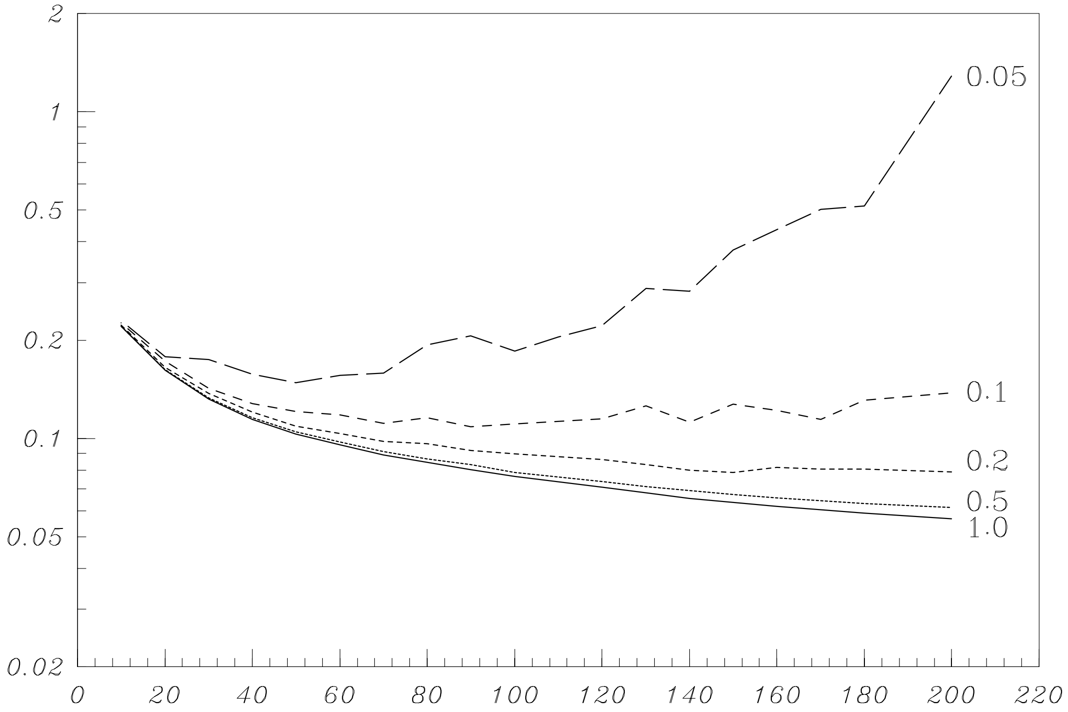

Figure 3: Mean squared error versus the number of fitting eigenfunctions for snapshots that were not part of the original ensemble. Different curves show the error for different unmasked areas, p. When $p = 1$ the entire picture is unmasked and the mean squared error is the best that may be attained for a given number of fitting eigenfunctions.

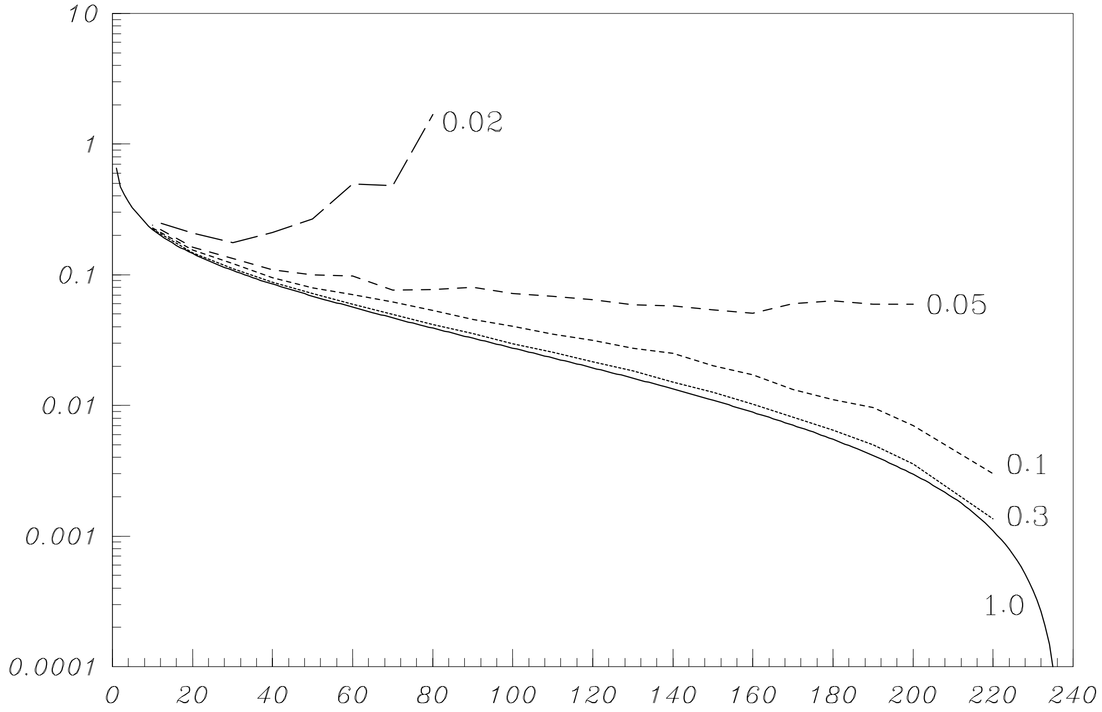

Figure 4: Mean squared error versus the number of fitting eigenfunctions for snapshots that were members of the original ensemble. Different curves show the error for different unmasked areas, p. When $p = 1$ the entire picture is unmasked and the mean squared error is the best that may be attained for a given number of fitting eigenfunctions.

Figure 5: Reconstruction of a monkey face using eigenfunctions derived from human faces. Left: unobscured face. Middle: reconstruction from a $p = 0.5$ random mask, using 100 eigenfunctions. Right: reconstruction using 220 eigenfunctions and the entire area.

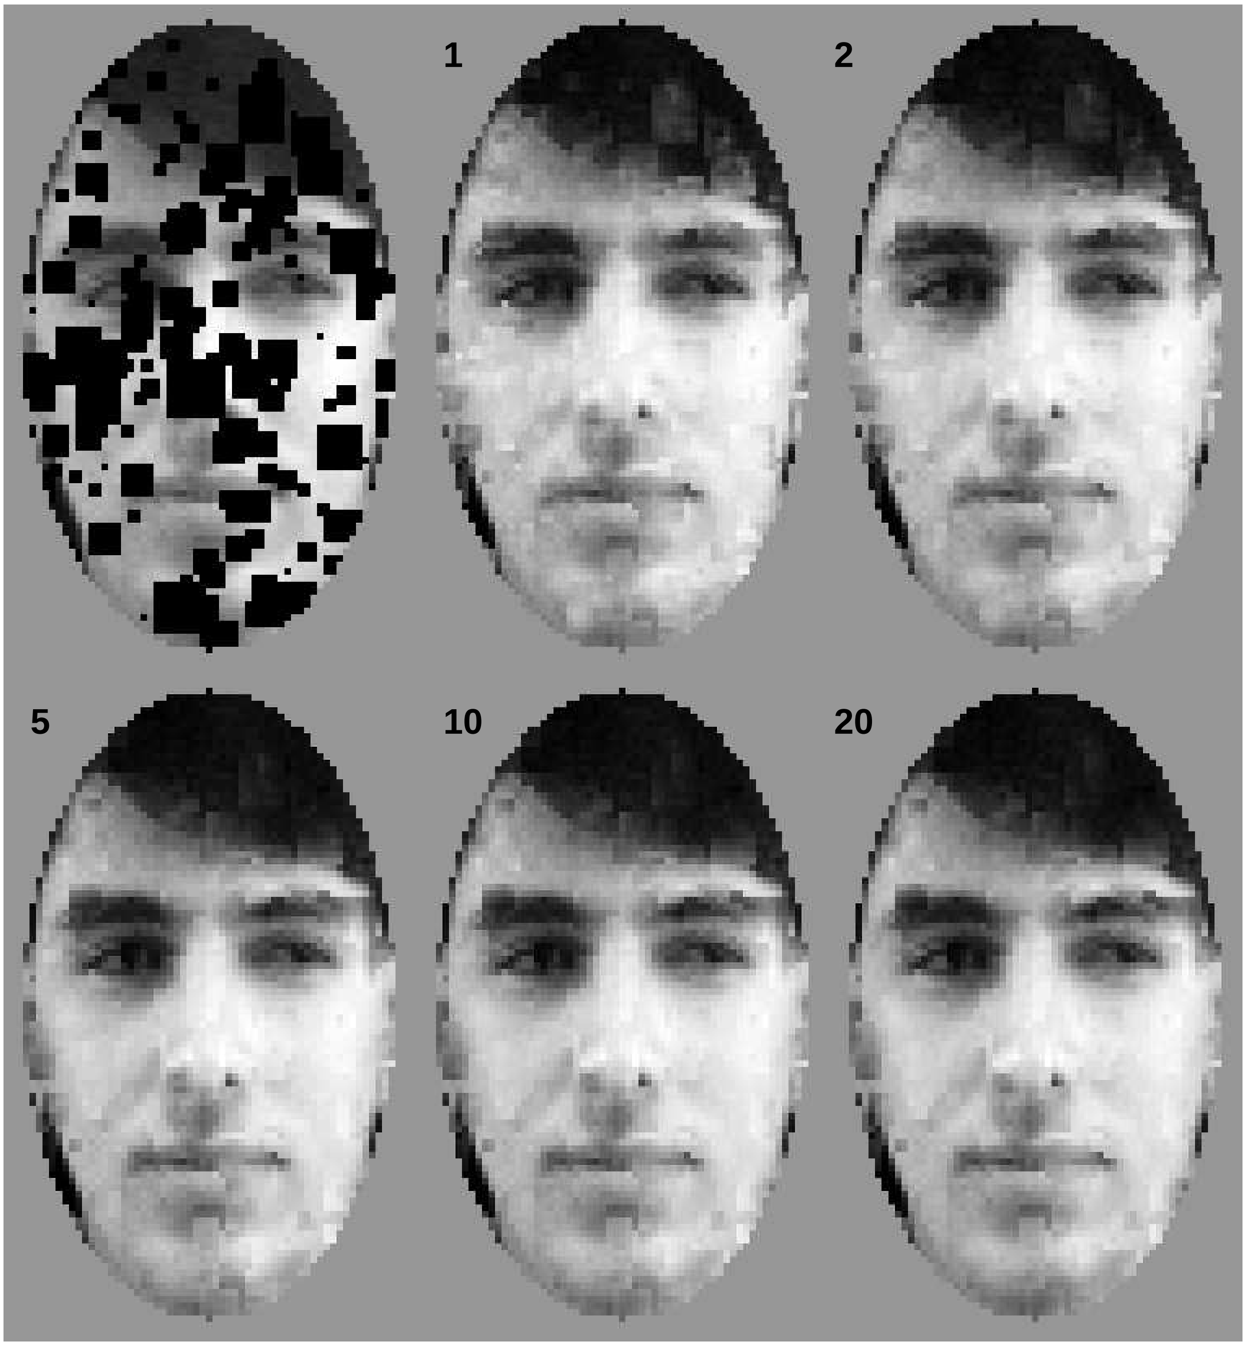

Figure 6: A 40% masked face $\tilde{\phi}_{53}$and the intermediate snapshots $\tilde{\phi}_{53}^{(1)}$, $\tilde{\phi}_{53}^{(2)}$, $\tilde{\phi}_{53}^{(5)}$,$\tilde{\phi}_{53}^{(10)}$ and $\tilde{\phi}_{53}^{(20)}$ as the iteration scheme proceeds. At each stage the masked regions have been repaired with $R = 30$ eigenfunctions derived from the snapshots from the previous iteration.

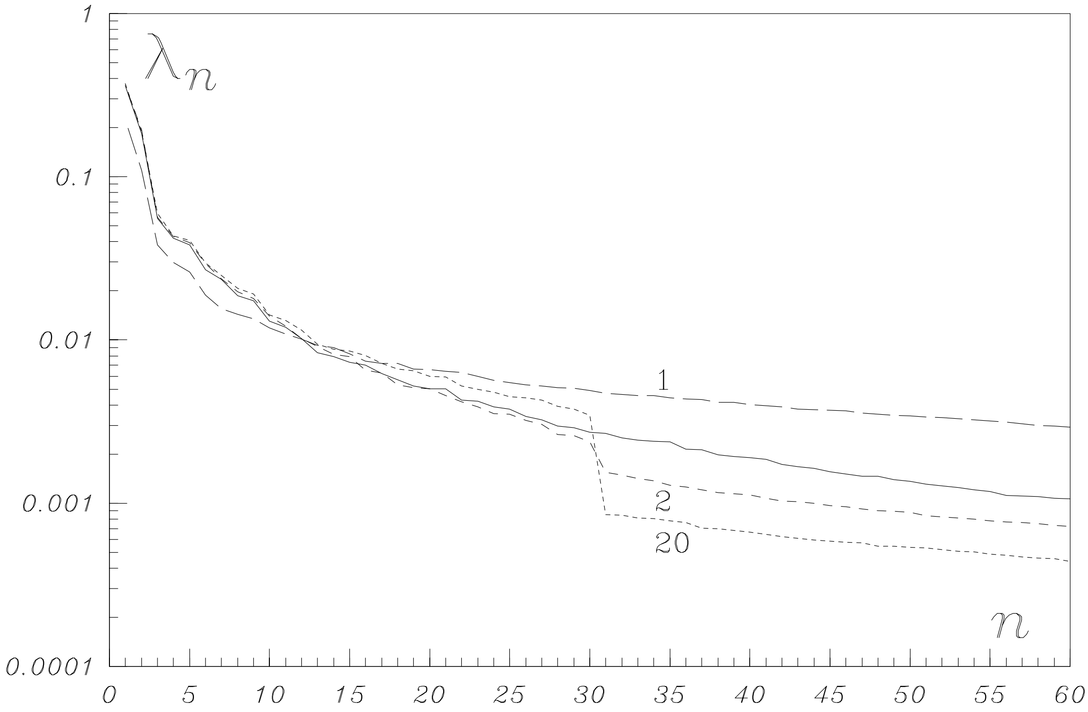

Figure 7: Eigenvalue spectrum, λ_n, after 1, 2 and 20 iterations. The solid line shows the spectrum derived from the unmarred snapshots.

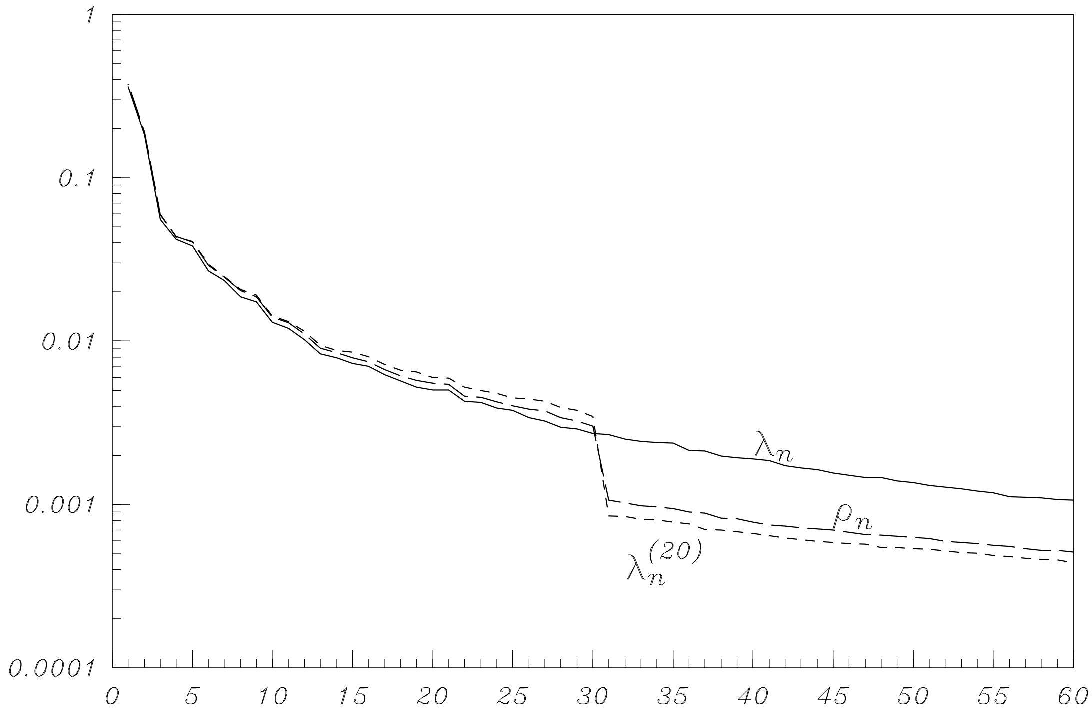

Figure 8: Eigenvalue spectrum, $\lambda_n^{(20)}$, after 20 iterations compared with the unmarred spectrum, λ_n, and the spectrum, ρ_n, derived from the snapshots repaired with eigenfunctions from the unmarred ensemble.

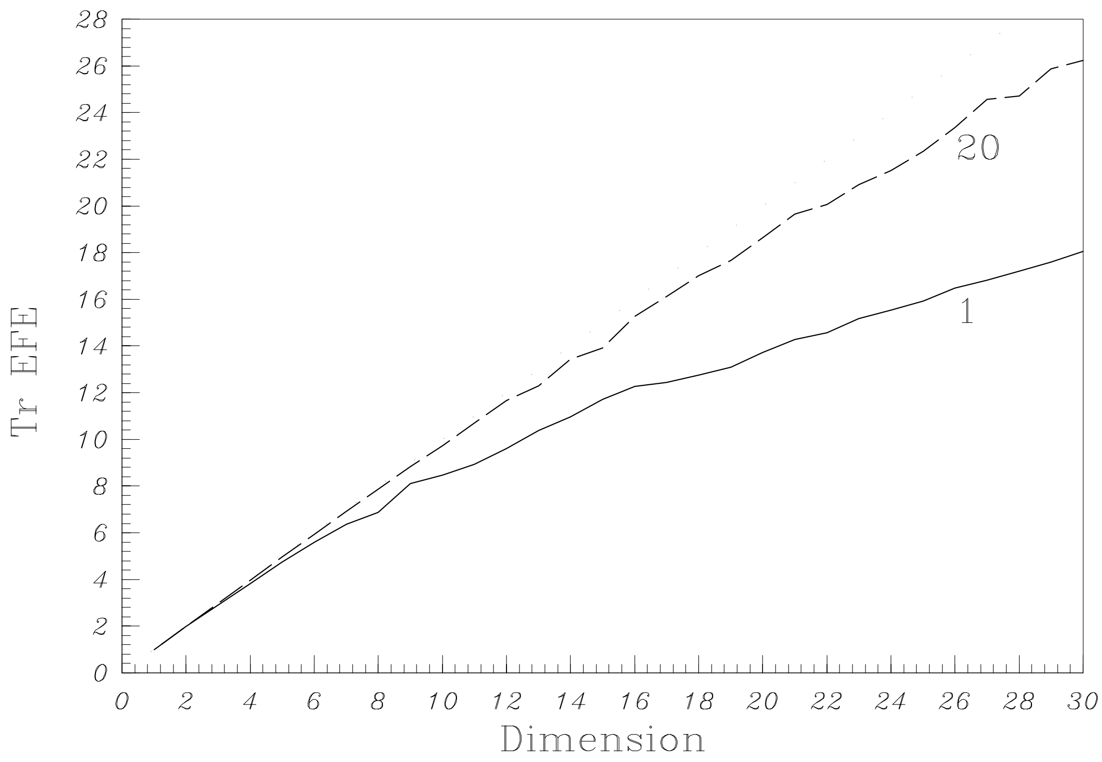

Figure 9: Commonality of subspaces spanned by the unmarred eigenfunctions, ψ_n, and the repaired eigenfunctions, ψ_n^k, after $k = 1$ and $k = 20$ iterations.

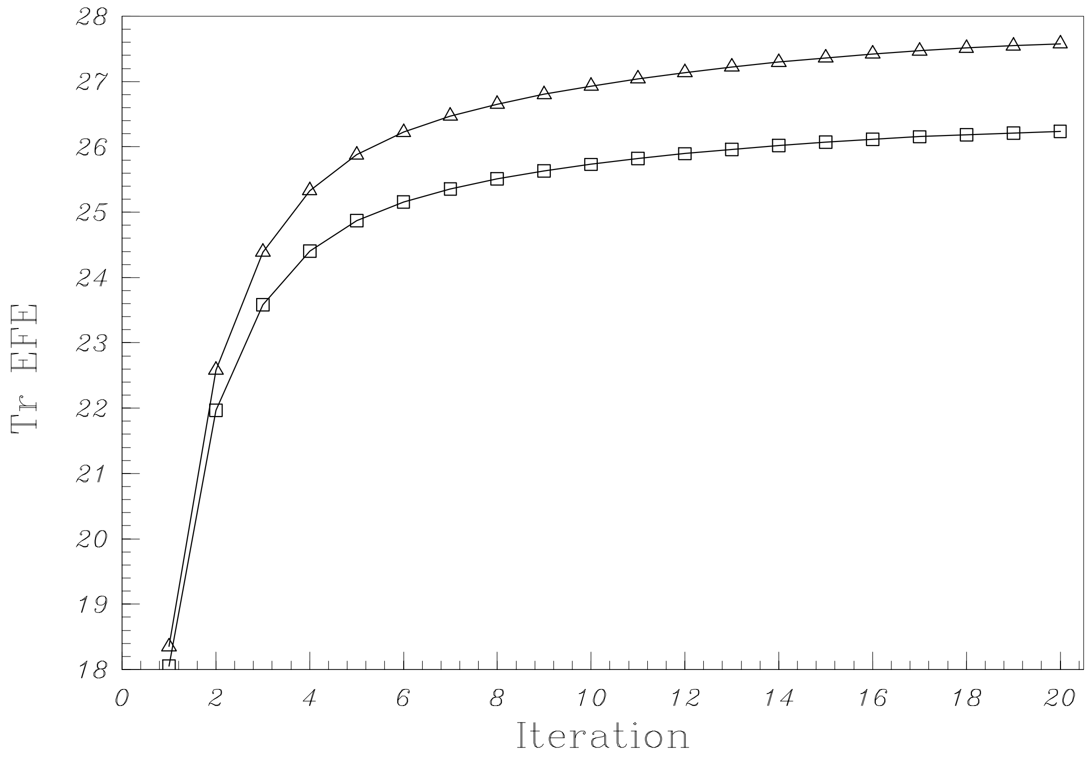

Figure 10: Convergence with iteration of the subspaces spanned by the repaired eigenfunctions and the unmarred eigenfunctions (squares) and the subspaces spanned by the repaired eigenfunctions and the eigenfunctions derived from marred data repaired with perfect eigenfunctions (triangles).

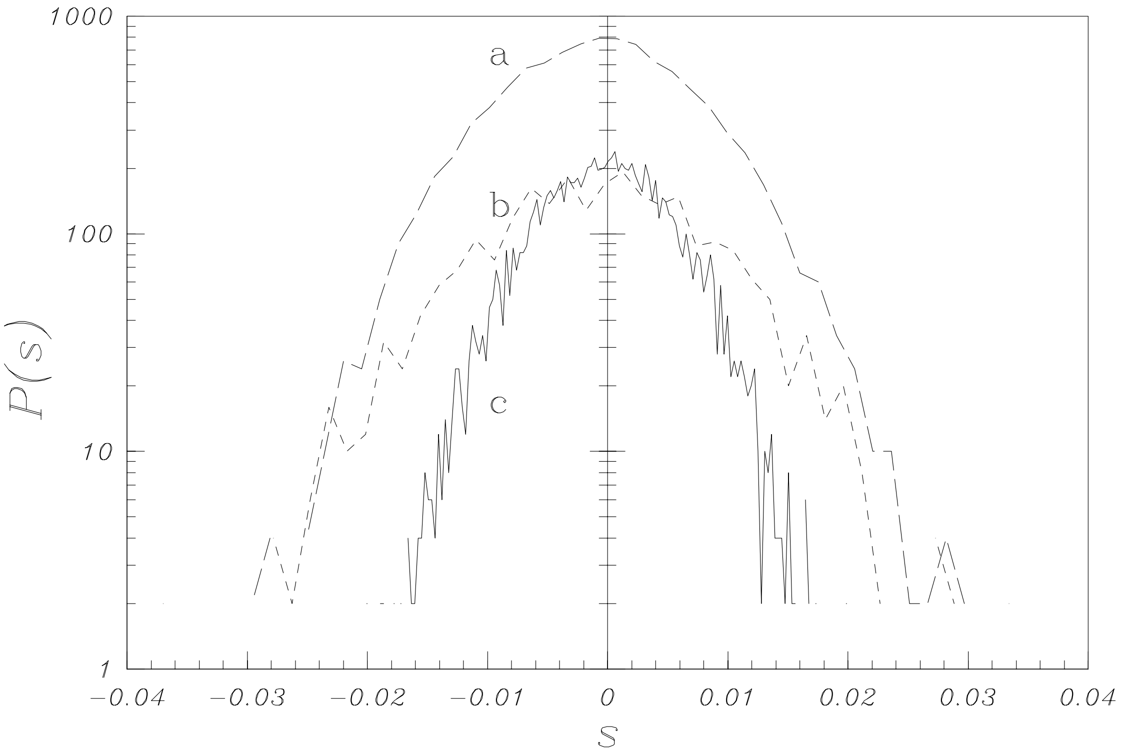

Figure 11: Empirical probability densities, $P(s)$, of the off-diagonal elements of the matrix $\mathbf{M}$, illustrating the Gaussian nature of the distribution. The matrices were generated at the following parameters: (a) $p = 0.7$, $N = 100$; (b) $p = 0.7$, $N = 50$; (c) $p = 0.1$, $N = 100$.