Biol. Cybern. 77, 407–417 (1997)

Biological Cybernetics

Separating spatially distributed response to stimulation from background. I. Optical imaging

R. Everson[1], **B.W. Knight**[2], **L. Sirovich**[1]

[1] Laboratory for Applied Mathematics, CUNY/Mount Sinai, New York, NY 10029, USA
[2] Rockefeller University, New York, NY 10021, USA

Received: 17 January 1997 / Accepted in revised form: 4 August 1997

Abstract. We consider the problem of estimating a small stimulus-induced response to stimulation that is masked by a fluctuating background when measurements of the background in the absence of stimulation are available, as is common in optical imaging of the cortex and in many other experimental situations. Two related methods based on the Karhunen-Loève procedure are discussed. One seeks the function, an indicator function, that is most parallel to the response data and most orthogonal to the background data. The second removes the subspace spanned by the background from the response. Numerical investigations on simulated optical imaging data show that the first method is generally superior. Connections between the two methods and techniques for assessing the quality of the result are discussed.

1 Introduction

In a number of applications, one encounters the challenge of extracting a very small signal submerged in a dominating background. This is the case that exists, for example, for optical imaging of the mammalian visual cortex (Blasdel and Salama 1986; Ts'o et al. 1990; Frostig et al. 1990) in response to external visual stimulation. In this instance, which will be used to illustrate the methods presented in this paper, the signal amplitude is between 10^{-3} and 10^{-4} of the background "noise". The extremely small changes in cortical reflectance, elicited in response to the stimulus, sit in a background that fluctuates due to respiration and heart beat, autonomous neural activity (Arieli et al. 1996), and other unwanted signals such as instrumental random noise.

Under ordinary conditions extraction of a meaningful signal from such noisy records might pose near-insurmountable difficulties. While the task is quite formidable, the challenge is made less daunting by the fact that the nature and timing of the stimulation are known a priori, as are some of the sources of the noise. This separation of the data into classes lies at the heart of the procedures that we present below. Examples of the application of this technique to optical imaging are given in Sirovich et al. (1996) and Everson et al. (1997 and in preparation).

Correspondence to: R. Everson (e-mail: rme@camelot.mssm.edu)

2 Background

The data which we analyze are typically in the form of images and, although the procedures hold for data of more general format, it is convenient to refer metaphorically to the data as composed of images. Thus we denote the data set by

$$f = f(t, \mathbf{x}) \tag{1}$$

where t is the index or time-stamp, and f measures the 'gray level' at 'pixel' location $\mathbf{x}$ of the image collected at time t.

An optimal organization of the image information is provided by the Karhunen-Loève (K-L) procedure (Karhunen 1946; Loève 1955; Stewart 1993; Schmidt 1907; Sirovich and Everson 1992). Since this formalism is a key element in our later discussion, we briefly outline the procedure. To begin, we attempt to expand the data set (1) in the form

$$f(t, \mathbf{x}) = \sum_n a_n(t)\sigma_n\psi_n(\mathbf{x}) \tag{2}$$

where both $\{a_n\}$ and $\{\psi_n\}$ form orthonormal sets. If (2) is to hold under these orthonormality conditions it follows that

$$\sigma_n\psi_n(\mathbf{x}) = \sum_t a_n(t)f(t, \mathbf{x}) = (a_n, f)_t \tag{3}$$

where $(f, g)_t$ denotes the inner product with respect to the temporal variable, t, and

$$\sigma_n a_n(t) = \sum_{\mathbf{x}} \psi(\mathbf{x})f(t, \mathbf{x}) = (\psi_n, f)_{\mathbf{x}} \tag{4}$$

where $(f, g)_{\mathbf{x}}$ denotes the inner product with respect to the spatial variable, $\mathbf{x}$. If (4) is back-substituted into (3) we obtain

$$\sum_s C(t, s)a_n(s) = \lambda_n a_n(t) \tag{5}$$

where $\lambda_n = \sigma_n^2$, and

$$C(t,s) = \sum_{\mathbf{x}} f(t,\mathbf{x}) f(s,\mathbf{x}) \tag{6}$$

is proportional to the temporal correlation. Alternatively, if (3) is substituted in (4) we obtain

$$\sum_{\mathbf{y}} K(\mathbf{x},\mathbf{y})\psi_n(\mathbf{y}) = \lambda_n \psi_n(\mathbf{x}) \tag{7}$$

where

$$K(\mathbf{x},\mathbf{y}) = \sum_{t} f(t,\mathbf{x}) f(t,\mathbf{y}) \tag{8}$$

is proportional to the spatial covariance.

Both kernels are symmetric, $C(t,s) = C(s,t)$ and $K(\mathbf{x},\mathbf{y}) = K(\mathbf{y},\mathbf{x})$, from which it follows that the corresponding eigenfunctions form complete orthonormal sets. Thus the existence of the expansion (2) is established. Equations (3) and (4) state that if one eigenset is known, the other is determined by quadrature. From a practical point of view one solves (5) or (7) depending on whether the number of images is greater than the number of pixels in each image. When the number of snapshots is smaller the eigenvectors are more conveniently determined via (5), and the technique is known as the *method of snapshots* (Sirovich 1987).

The eigenfunctions $\psi_n(\mathbf{x})$ for the spatial domain are known variously as the Karhunen-Loève eigenfunctions, empirical eigenfunctions or empirical orthogonal functions (Lorenz 1956) and they may be identified with the principal components of principal components analysis. It is convenient to arrange discrete data as a matrix $\mathbf{F}$, each of whose rows is an image ($F_{nm} = f(t_n, \mathbf{x}_m)$). In this form the K-L decomposition may be recognized as a singular value decomposition of $\mathbf{F}$:

$$\mathbf{F} = \mathbf{A}\boldsymbol{\Sigma}\boldsymbol{\Psi}^{\mathrm{T}} \tag{9}$$

where $\mathbf{A}$ is an orthogonal matrix whose columns $\{a_n(t)\}$ span the range of $\mathbf{F}$, $\boldsymbol{\Sigma}$ is the diagonal matrix whose entries are σ_n and $\boldsymbol{\Psi}$ is an orthogonal matrix whose columns are the K-L eigenvectors $\{\psi(\mathbf{x})\}$. Thus (9) is (2) written in vector notation.

The eigenvalues, which are non-negative, measure the mean square projection of each eigenfunction onto the data

$$\lambda_n = \sigma_n^2 = \sum_t (\psi_n, f)_{\mathbf{x}}^2 = \sum_{\mathbf{x}} (a_n, f)_t^2 \tag{10}$$

and the data may be approximated by projecting onto those eigenfunctions (the 'principal eigenfunctions') which carry significant power. In fact the K-L decomposition provides the optimal basis, among all linear bases, for such approximation since the mean squared error,

$$\| f(t,\mathbf{x}) - \sum_{n=1}^{N} a'_n(t)\sigma'_n\psi'_n(\mathbf{x}) \|_{\mathbf{x},t}^2 \tag{11}$$

is minimized for any choice of N by the choice $a'_n = a_n$, $\sigma'_n = \sigma_n$ and $\psi'_n(\mathbf{x}) = \psi_n(\mathbf{x})$.

3 Indicator functions

The images which are of interest have an extremely low signal-to-noise ratio. However, since they carry responses to known stimuli, presented at known times, we can use this information to extract the small signal.

In a wide set of circumstances the image data can be naturally divided into two classes, which we will denote by $f(t,\mathbf{x})$ and $\hat{f}(\hat{t},\mathbf{x})$. For example, f can represent the ensemble of images collected in response to a specific stimulus, while $\hat{f}$ are a set of reference images, the 'response' images in the absence of a stimulus – what we refer to as blanks, or the class of all other stimuli. In what follows we take t to be a discrete index, i.e., we consider $f(t,\mathbf{x})$, $t = 1, 2, \ldots, N$; and $\hat{f}(\hat{t},\mathbf{x})$, $\hat{t} = 1, 2, \ldots, M$. Throughout this article we refer to the data as images; however, it should be emphasized that the two-dimensional nature of the images is not explicitly exploited and $f(t,\mathbf{x})$ could be vectors, pixels, voxels or higher-dimensional data.

The standard 'differential imaging' procedure (Blasdel 1992) for extracting the signal is to subtract the average blank from the average response; namely

$$\begin{aligned} \Delta(\mathbf{x}) &= \langle f(t,\mathbf{x})\rangle - \langle \hat{f}(\hat{t},\mathbf{x})\rangle \\ &= \frac{1}{N}\sum_{t=1}^{N} f(t,\mathbf{x}) - \frac{1}{M}\sum_{\hat{t}=1}^{M} \hat{f}(\hat{t},\mathbf{x}) \end{aligned} \tag{12}$$

In the limit of large N and M, and noise that is uncorrelated with the signal, $\Delta(\mathbf{x})$ estimates the true signal, approaching it as $O(\min(M,N)^{1/2})$. However, when N and M are of limited size, background contributions to the averages in (12) may fail to cancel.

We next present two, somewhat related, methods for signal extraction. These will be compared in Sect. 4.

3.1 Method I

We seek an image, known as the indicator function (or indicator image), $\phi(\mathbf{x})$ with which we attempt to satisfy the two conditions:

$$\begin{aligned} (f(t,\mathbf{x}), \phi(\mathbf{x}))_{\mathbf{x}} &= 1, \quad t = 1, \ldots, N \\ (\hat{f}(\hat{t},\mathbf{x}), \phi(\mathbf{x}))_{\mathbf{x}} &= 0, \quad \hat{t} = 1, \ldots, M \end{aligned} \tag{13}$$

Thus, in loose terms, the indicator function, ϕ, is aligned with the ensemble carrying the signal, f, and orthogonal to the complementary ensemble, $\hat{f}$. This construction bears a kinship to the linear discriminants of Fisher, used for deciding to which of two classes a particular vector belongs (Fisher 1936; Duda and Hart 1973). If we are successful in finding a ϕ that satisfies (13) then $(\phi, f)_{\mathbf{x}}$ is a clear *indicator* of whether f belongs to the stimulated or to the blank class.

In usual practice the number of image pixels, $P = \dim[\phi] = \dim[f] = \dim[\hat{f}]$, greatly exceeds the number of images, i.e., $P > N + M$. It would appear, therefore, that the system is underdetermined. On the other hand it is not clear that the right-hand side of (13) lies in the range of

the left-hand side. The first possibility is confronted below, while the second can be circumvented if instead we seek to minimize the criterion function:

$$\mathscr{C}(\phi) = \sum_{t=1}^{N}((f,\phi)_{\mathbf{x}} - 1)^2 + \sum_{\hat{t}=1}^{M}(\hat{f},\phi)_{\mathbf{x}}^2 \tag{14}$$

It is straightforward to compute $\delta\mathscr{C}/\delta\phi$, and setting it to zero we obtain

$$\sum_{\mathbf{y}}\left(\sum_{t} f(t,\mathbf{x})f(t,\mathbf{y}) + \sum_{\hat{t}}\hat{f}(\hat{t},\mathbf{x})\hat{f}(\hat{t},\mathbf{y})\right)\phi(\mathbf{y}) = \sum_{t} f(t,\mathbf{x}) \tag{15}$$

The operator of the left-hand side is the covariance of a total ensemble of images, and the right-hand side is proportional to the conditional average, averaged over only members of the distinguished class. Thus, we might write (15) in the form

$$\sum_{\mathbf{y}} K(\mathbf{x},\mathbf{y})\phi(\mathbf{y}) = \bar{f}(\mathbf{x}) = N\langle f\rangle \tag{16}$$

Clearly, the solution to (16) involves the inversion of the operator K, which requires special attention since the eigenvalues of K generally decay slowly to zero. In fact, it is desirable to filter out pixel noise and contributions to the indicator function that result from chance differences between the stimulated and blank datasets (as opposed to the stimulus-driven component of ϕ, which is sought). To this end ϕ is restricted to lie in the space spanned by the first T eigenfunctions of $K(\mathbf{x},\mathbf{y})$; thus if

$$K_T(\mathbf{x},\mathbf{y}) = \sum_{n=1}^{T}\sigma_n^2\psi_n(\mathbf{x})\psi_n(\mathbf{y}) \tag{17}$$

the indicator function is given by the solution of

$$\sum_{\mathbf{y}} K_T(\mathbf{x},\mathbf{y})\phi(\mathbf{y}) = N\langle f\rangle \tag{18}$$

The choice of the truncation T, which requires some care, is addressed in Sect. 5.

The snapshot method reduces the problem by seeking a solution, ϕ, of (15) in the space of images

$$\phi = \sum_{t=1}^{N} u(t)f(t,\mathbf{x}) + \sum_{\hat{t}=1}^{M} v(\hat{t})\hat{f}(\hat{t},\mathbf{x}) \tag{19}$$

The variational form of (14) with respect to $u(t)$ and $v(\hat{t})$ is satisfied if the following equation in terms of the temporal correlations is true:

$$\begin{bmatrix}\mathbf{C} & \mathbf{B}\\ \mathbf{B}^T & \hat{\mathbf{C}}\end{bmatrix}\begin{bmatrix}\mathbf{u}\\ \mathbf{v}\end{bmatrix} = \begin{bmatrix}\mathbf{1}_N\\ \mathbf{0}_M\end{bmatrix} \tag{20}$$

where

$$\begin{aligned} C_{nm} &= (f(n,\mathbf{x}), f(m,\mathbf{x}))_{\mathbf{x}}\\ B_{pq} &= (f(p,\mathbf{x}), \hat{f}(q,\mathbf{x}))_{\mathbf{x}}\\ \hat{C}_{kl} &= (\hat{f}(k,\mathbf{x}), \hat{f}(l,\mathbf{x}))_{\mathbf{x}} \end{aligned} \tag{21}$$

3.2 Method II

Two features of the procedure presented in the previous section deserve further consideration: (i) under experimental conditions, considerable variation in the members of $f(t,\mathbf{x})$ can be expected, even though the images are obtained under like conditions; (ii) in general, the response to stimulation may be multi-dimensional, so we might expect more than one direction for which (13) holds. We therefore consider the following variational problem for a (different) indicator function, $\phi'(\mathbf{x})$:

$$\begin{aligned} &\text{extremize} \quad \sum_{t}(f,\phi')_{\mathbf{x}}^2\\ &\text{subject to} \quad \|\phi'\|_{\mathbf{x}}^2 = 1, \text{ and } (\hat{f},\phi')_{\mathbf{x}} = 0 \text{ for } \hat{t} = 1,\ldots,M \end{aligned} \tag{22}$$

To carry this out we introduce Lagrange multipliers γ and $\mu(\hat{t})$, and the criterion function

$$\mathscr{C} = \sum_{t}(f,\phi')_{\mathbf{x}}^2 - \gamma(\phi',\phi')_{\mathbf{x}} - \sum_{t}\mu(t)(\hat{f},\phi')_{\mathbf{x}} \tag{23}$$

On setting the variation of $\mathscr{C}$ with respect to ϕ' to zero we obtain,

$$\sum_{\mathbf{y}}\sum_{t}(f(t,\mathbf{x}), f(t,\mathbf{y}))_t\phi'(\mathbf{y}) = \gamma\phi'(\mathbf{x}) + \sum_{\hat{t}}\mu(\hat{t})\hat{f}(\hat{t},\mathbf{x}) \tag{24}$$

We solve as above by taking

$$\phi' = \sum_{t=1}^{N}\alpha(t)f(t,\mathbf{x}) + \sum_{\hat{t}=1}^{M}\beta(\hat{t})\hat{f}(\hat{t},\mathbf{x}) \tag{25}$$

If this is substituted into (24) we obtain

$$\sum_{k=1}^{N} C_{nk}\alpha_k + \sum_{l=1}^{M} B_{nl}\beta_l = \gamma\alpha_n \tag{26}$$

$$\mu_m + \gamma\beta_m = 0 \tag{27}$$

Note that the second of these only serves to evaluate β_m. The system is completed by applying the side condition $(\hat{f},\phi')_{\mathbf{x}} = 0$ to (25). This yields

$$\sum_{l=1}^{M} C_{nl}\beta_l + \sum_{k=1}^{N}\alpha_k B_{kn} = 0 \tag{28}$$

Therefore the coefficients are determined by

$$\begin{bmatrix}\mathbf{C} & \mathbf{B}\\ \mathbf{B}^{\mathrm{T}} & \hat{\mathbf{C}}\end{bmatrix}\begin{bmatrix}\boldsymbol{\alpha}\\ \boldsymbol{\beta}\end{bmatrix} = \gamma\begin{bmatrix}\boldsymbol{\alpha}\\ \mathbf{0}\end{bmatrix} \tag{29}$$

or equivalently

$$\begin{bmatrix}\mathbf{C}-\gamma & \mathbf{B}\\ \mathbf{B}^{\mathrm{T}} & \hat{\mathbf{C}}\end{bmatrix}\begin{bmatrix}\boldsymbol{\alpha}\\ \boldsymbol{\beta}\end{bmatrix} = \mathbf{0} \tag{30}$$

The last relation shows that, in general, we obtain N real non-negative eigenvalues, γ_n, $n = 1,\ldots,N$. Then if we write the corresponding eigenvectors as $[\boldsymbol{\alpha}_n,\boldsymbol{\beta}_n]^{\mathrm{T}}$, it follows that $(\boldsymbol{\alpha}_n,\boldsymbol{\alpha}_m) = 0$ for distinct eigenvalues γ_m and γ_n.

We may now relate the indicator functions of Methods I and II. Without loss of generality we can take $\|\boldsymbol{\alpha}_n\|^2 = 1$. Then it follows that $[\mathbf{u}, \mathbf{v}]^{\mathrm{T}}$ of Method I is given by

$$\begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \sum_{n=1}^{N} \frac{(\boldsymbol{\alpha}_n, \mathbf{1})}{\gamma_n} \begin{bmatrix} \boldsymbol{\alpha}_n \\ \boldsymbol{\beta}_n \end{bmatrix} \tag{31}$$

The derivation of this relation assumes that the Method I indicator function is not restricted to a truncated subspace. Nonetheless, the results presented in Sect. 4 show that the eigenvector of (30) corresponding to the largest eigenvalue is similar to the Method I indicator function.

The nature of the indicator functions is illuminated by noting that if the spaces spanned by the stimulated and blank ensembles do not intersect (so that $B_{pq} = (f(p, \mathbf{x}), \hat{f}(q, \mathbf{x}))_{\mathbf{x}} = 0$) the $\boldsymbol{\alpha}_n$ are eigenvectors of $\mathbf{C}$, the temporal correlation matrix of the stimulated ensemble. The indicator functions ϕ' are therefore identified with the K-L eigenvectors of the stimulated ensemble. When $\mathbf{B}$ is not zero the $\boldsymbol{\alpha}_n$ are eigenfunctions of the operator

$$\mathbf{Q} = \mathbf{C} - \mathbf{B}\hat{\mathbf{C}}^{-1}\mathbf{B}^{\mathrm{T}} \tag{32}$$

Writing $\mathbf{Q}$ in terms of $\mathbf{F}$ and the singular value decomposition (9) of $\hat{\mathbf{F}} = \hat{\mathbf{A}}\hat{\boldsymbol{\Sigma}}\hat{\boldsymbol{\Psi}}^{\mathrm{T}}$, obtains

$$\mathbf{Q} = \mathbf{F}\mathbf{F}^{\mathrm{T}} - (\mathbf{F}\,\hat{\boldsymbol{\Psi}}\hat{\boldsymbol{\Sigma}}\hat{\mathbf{A}}^{\mathrm{T}})(\hat{\mathbf{A}}\hat{\boldsymbol{\Sigma}}^{-2}\hat{\mathbf{A}}^{\mathrm{T}})(\hat{\mathbf{A}}\hat{\boldsymbol{\Sigma}}\hat{\boldsymbol{\Psi}}^{\mathrm{T}}\,\mathbf{F}^{\mathrm{T}}) \tag{33}$$

$$= \mathbf{F}\mathbf{F}^{\mathrm{T}} - \mathbf{F}\hat{\boldsymbol{\Psi}}\hat{\boldsymbol{\Psi}}^{\mathrm{T}}\mathbf{F}^{\mathrm{T}} \tag{34}$$

$$= (\mathbf{F} - \mathbf{F}\hat{\mathbf{P}})(\mathbf{F} - \mathbf{F}\hat{\mathbf{P}})^{\mathrm{T}} \tag{35}$$

where $\hat{\mathbf{P}} = \hat{\boldsymbol{\Psi}}\hat{\boldsymbol{\Psi}}^{\mathrm{T}}$ is an orthogonal projector onto the space spanned by the blank ensemble. $\mathbf{Q}$ is thus recognized as the temporal correlation matrix for the stimulated ensemble from which the subspace of blanks has been removed. The indicator functions ϕ' are therefore the K-L eigenfunctions of this reduced ensemble. Here lies the essential difference between the two methods: the ϕ' are forced to be strictly orthogonal to every element in the blank space, while the demand of orthogonality of ϕ to the blank space is only satisfied in an average sense (cf. (14)).

Rather than remove the entire blank space from the stimulated space it may be desirable to remove only part of it. We can modify the definition of ϕ' so that it is orthogonal only to T of the K-L eigenfunctions $\hat{\psi}_n$ of the blank space; that is, ϕ' is the solution to the following variational problem:

$$\text{extremize} \quad \sum_{n=1}^{N} (f, \phi')_{\mathbf{x}}^2$$

$$\text{subject to} \quad \|\phi'(\mathbf{x})\|^2 = 1 \text{ and } (f, \hat{\psi}_n)_{\mathbf{x}} = 0, \text{ for } n = 1, \ldots, T \tag{36}$$

This endows the type II indicator image with a noise-cut-off parameter somewhat similar to that introduced for type I.

4 Numerical examples

In order to illustrate the indicator functions and to investigate their efficacy we have studied their application to real optical imaging data onto which a synthetic response was added. The unstimulated or blank data were composed of 400 optical images, recorded in the absence of visual stimulation, from the striate cortex of macaque (for details see O'Brien et al. 1995). The response to stimulation was modeled by adding a function $\nu(t)p(\mathbf{x})$ to a further 400 images that were recorded interleaved with the blank data. The spatial pattern (Fig. 1a) [in coordinates $\mathbf{x} = (x, y)$, $-64 \le x < 64$, $-48 \le y < 48$, and with $\mathbf{k} = (10, 20)$] was

$$p(x, y) = \begin{cases} \frac{ReD}{|D|} \text{ with } D(x, y) = (x + iy)e^{i\mathbf{k}\cdot\mathbf{x}} & x < 2y - 32 \\ 0 & \text{elsewhere} \end{cases} \tag{37}$$

which is a rough caricature of ocular dominance columns found in the primate visual cortex. The amplitudes $\nu(t)$ were chosen from pseudo-random numbers uniformly distributed between 0 and 1, so that $\nu(t)p(\mathbf{x})$ is $O(1)$, which is comparable with the amplitude of typical responses in optical imaging data. Stimulated and blank images are indistinguishable to the naked eye (Fig. 1). The mean pixel value, $\frac{1}{P}\sum_{\mathbf{x}}^{P} \bar{f}(\mathbf{x}) = \frac{1}{P}\sum_{\mathbf{x}}^{P}\langle f(t, \mathbf{x})\rangle$, of the simulated and blank data together was 1975. The root mean square fluctuation about the mean image, $\langle \frac{1}{P}\sum_{\mathbf{x}}^{P}(f(t, \mathbf{x}) - \bar{f}(\mathbf{x}))^2\rangle^{1/2}$ was 37.35. Also shown in Fig. 1d is the subtraction image $\Delta(\mathbf{x})$, in which it is hard to discern p. Of course, in this example a simple spatial Fourier filter, aligned with the pattern, would do an excellent job of extracting, p. However, in real applications the sought spatial pattern is unknown, precluding the design of a filter to extract it.

The spectrum of the eigenvalues, λ_n, from a K-L decomposition of the combined stimulated and blank ensembles is shown in Fig. 2. The first spatial eigenfunction, $\psi_1(\mathbf{x})$, corresponding to the dominant eigenvalue lies close to the mean image. Eigenfunctions up to the knee at $n \approx 150$ describe fluctuations in the vegetative response of the cortex and the added $p(\mathbf{x})$ of (37). Beyond the knee the eigenfunctions describe mainly pixel noise and display only weak pixel correlations. Removal of pixel noise and computational economies are afforded by projecting the data onto only the earlier eigenfunctions and by performing subsequent calculations in this smaller subspace. The calculations presented here were performed in the space spanned by the first 500 eigenfunctions.

The degree to which the stimulated and blank subspaces intersect is quantified by

$$\text{Tr } P_n\hat{P}_n = \text{Tr } \hat{P}_n P_n \hat{P}_n = \text{Tr } (P_n\hat{P}_n)(P_n\hat{P}_n)^{\mathrm{T}} = |P_n\hat{P}_n|^2 \tag{38}$$

where P_n and $\hat{P}_n$ are projectors onto the principal n K-L eigenfunctions of the stimulated and blank spaces respectively. Here the operator $P_n\hat{P}_nP_n$ has eigenvectors in the stimulated n-dimensional subspace whose eigenvalues state how much survives after such a vector is projected into the unstimulated space and then back into the stimulated space. The trace is the sum of these eigenvalues. If the subspaces are identical Tr $P_n\hat{P}_n = n$, whereas if they do not intersect Tr $P_n\hat{P}_n = 0$. As shown in Fig. 3 the stimulated and blank

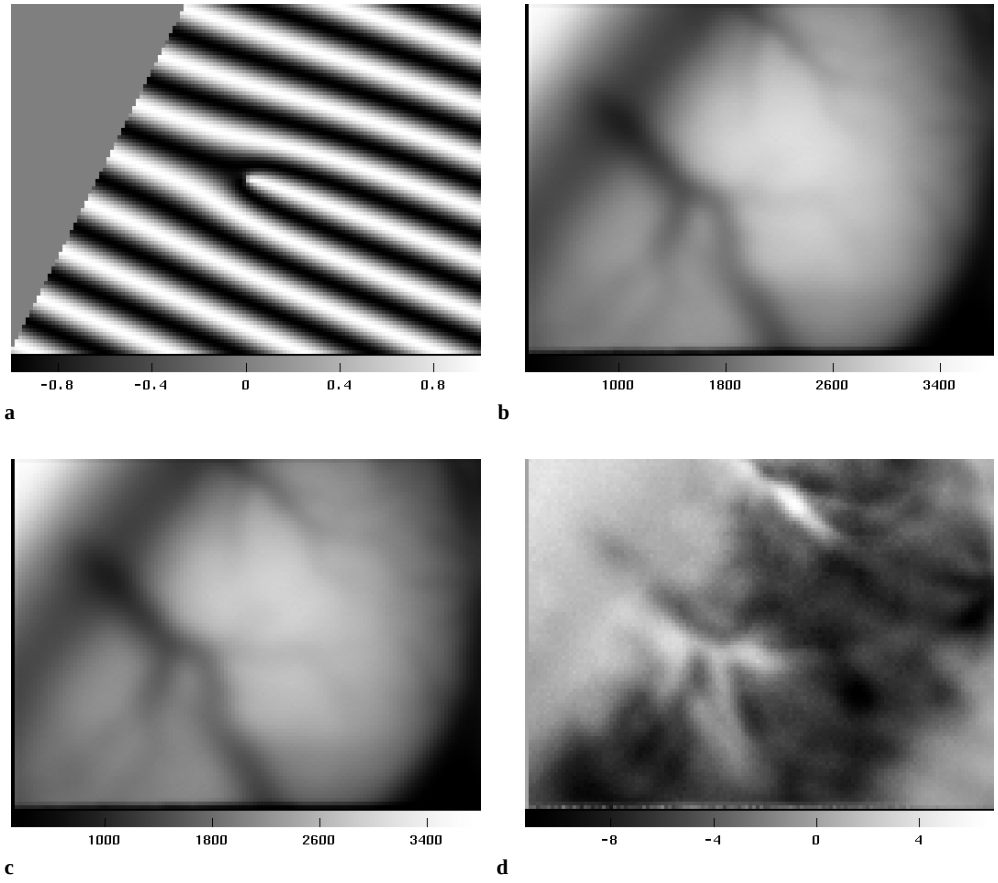

Fig. 1. a The pattern added to blank images. **b** A typical blank image. **c** A weighted sum used in the simulation, with **a** added at maximum weighting strength. **d** The difference image, (12), between the stimulated and blank ensembles

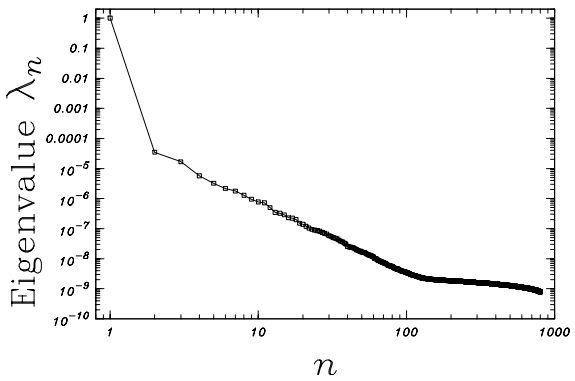

Fig. 2. Spectrum of eigenvalues from a Karhunen-Loève decomposition of the stimulated and blank data

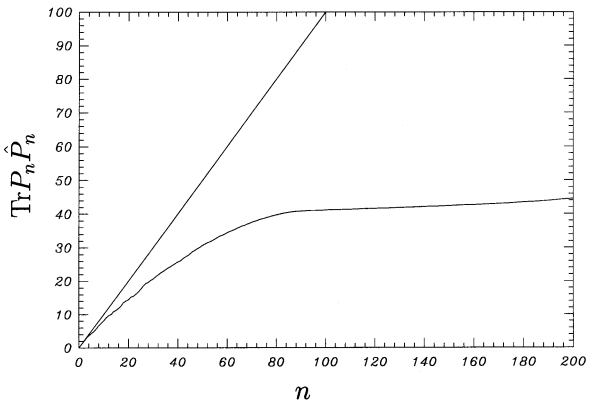

Fig. 3. Overlap between the stimulated and blank subspaces

subspaces differ considerably in the first 100 dimensions, in addition to the one-dimensional difference that may be expected from the presence of p(**x**) in the stimulated data. The effect of pixel noise is apparent for $n \geq 100$. Two sets, each of 400 samples of pixel noise, will have little in common, since in the whole picture there are 12 288 pixels, each independently affected by pixel noise.

If the amplitude of p were larger, its contribution to the dataset would be approximately represented by a linear combination of one or two eigenfunctions. In such circumstances signal extraction is a simple matter; however, in this case, as is typical, $(p, \psi_n)_{\mathbf{x}}$ is not negligible for a range of n.

4.1 Method I

Figure 4 shows the angle ($\cos\theta = (p, \phi)_{\mathbf{x}} / \|p\| \|\phi\|$) between the indicator function ϕ and the pattern p as the truncation point T (18) is increased to a point where the empirical eigenfunctions are dominated by pixel noise. At $T \approx 45$ the indicator function most closely approximates p; the angle of 39° corresponds to capturing $\cos^2 39° = 61\%$ of the power of p. This indicator function (Fig. 5) shows the principal features, including the dislocation, of p, and we point out that ϕ(**x**) is close to zero in the region $y > 2x$ where $p = 0$.

The initial empirical eigenfunctions of both the stimulated and blank images span very similar spaces (Fig. 3)

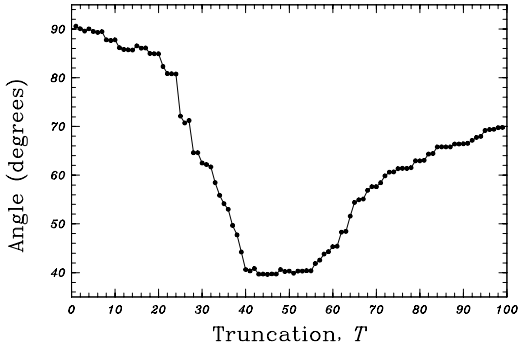

Fig. 4. Angle between the indicator function, ϕ, and the pattern p as a function of the truncation point, T

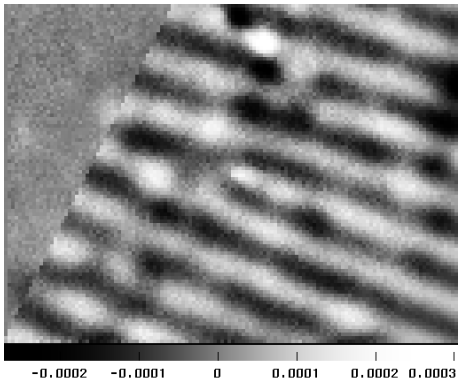

Fig. 5. Indicator function, $\phi(\mathbf{x})$, closest to the pattern p, calculated with truncation $T = 44$

and describe gross features of the data. They are practically orthogonal to p so that indicator functions constructed with small T fail to locate p. On the other hand, when T is large the indicator function describes not only p, but also differences between the datasets due to random fluctuations, and eventually pixel noise dominates the indicator function. Methods for determining the best T are the topic of Sect. 5.

4.2 Method II

Figure 6 shows the angle between p and the indicator function ϕ' (36) as the number of blank-space eigenfunctions that are removed is increased. After the first approximately 100 blank-space eigenfunctions have been removed only small decreases in angle are accrued as additional eigenfunctions, which describe mainly pixel noise, are subtracted. Figure 7 shows indicator functions corresponding to the removal of 90 and 400 eigenfunctions: there is little improvement in the angle at the expense of increased pixellation. The entire space spanned by the blanks does contain most of p, so removing sufficient blank eigenfunctions eventually degrades the quality of ϕ'.

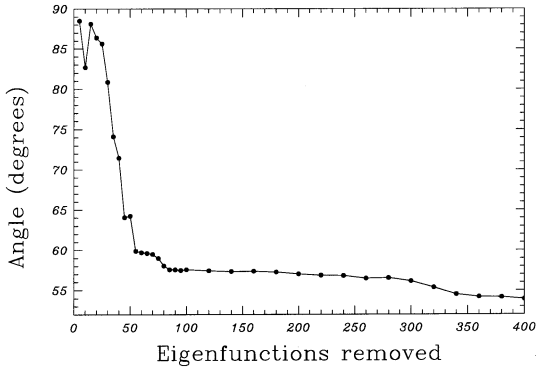

Fig. 6. Angle between the indicator function of Method II, ϕ', and the pattern p as a function of the number of blank-space eigenfunctions that are removed from the stimulated space

Although ϕ' in this case fails to locate p as accurately as ϕ, when the amplitude of p is larger or the blank and stimulated subspaces are more similar, the Method II indicator function can perform slightly better than ϕ. On the other hand, we have found that for real experimental data the Method I indicator function performs better; as discussed below, we suspect the reason to be that for real data the blank space is partially contaminated with pattern.

4.3 Contaminated data

The simulated data considered above are relatively clean in the sense that p does not lie in the span of the principal blank-space K-L eigenfunctions. Denoting the orthogonal projector onto the space spanned by the first n K-L eigenfunctions of the stimulated and blank spaces by P_n and $\hat{P}_n$ respectively, we have $\|P_{100}p\|/\|p\| = 0.98$ and $\|\hat{P}_{100}p\|/\|p\| = 0.26$. In many optical imaging experiments, however, the blank space is partially contaminated with p. To simulate this we add $\hat{\nu}(t)p(\mathbf{x})$ to the blank images, where $\langle\hat{\nu}(t)\rangle/\langle\nu(t)\rangle = 0.2$ so that the power carried by the pattern in the 'blank space' is 0.04 of that in the 'stimulated space'. In this case $\|\hat{P}_{100}p\|/\|p\| = 0.48$.

Figure 8 shows the angle between p and the two types of indicator function for the contaminated data. Here Method II throws out the baby with the bath-water because the pattern is spanned by the blank space and is therefore removed in the construction of ϕ'. It is our experience with real experimental data that Method. I indicator functions are superior to those of Method II. In the remainder of this paper we shall therefore concentrate on Method I.

5 Choosing the truncation

Construction of the indicator function formally involves inversion of the covariance operator $K(\mathbf{x}, \mathbf{y})$. As discussed in Sect. 3.1, the indicator function is restricted to lie in the space spanned by the first T eigenfunctions of $K(\mathbf{x}, \mathbf{y})$. The choice of the truncation T, however, requires some care. As Fig. 4 shows, if T is too small eigenfunctions of K spanning p are not included in the solution, whereas if T is too large contributions from chance differences between the stimulated and blank datasets dominate the solution.

In this section we discuss criteria for choosing the truncation. It is helpful to have an expression for the indicator functions as superpositions of K-L eigenfunctions. Substitution for K_T and f in (18) yields

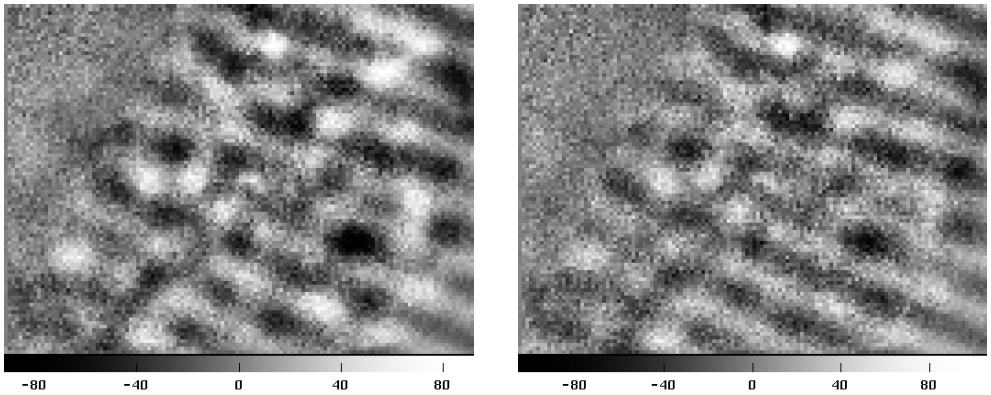

Fig. 7. Method II indicator functions, ϕ', corresponding to the removal of **a** 90 and **b** 400 blank-space eigenfunctions

$$\sum_{\mathbf{y}} \sum_{n=1}^{T} \sigma_n^2 \psi_n(\mathbf{x})\psi_n(\mathbf{y})\phi(\mathbf{y}) = \sum_{t=1}^{N} \sum_{k=1}^{T} a_k(t)\sigma_k \psi_k(\mathbf{x}) \quad (39)$$

where the sum on t extends over just the stimulated snapshots. Formation of the spatial inner product with both sides of (39) obtains

$$\sigma_n^2(\psi_m(\mathbf{y}), \phi(\mathbf{y}))_{\mathbf{x}} = \sum_{t} \sum_{k=1}^{T} a_k(t)\sigma_k \sum_{\mathbf{x}} \psi_k(\mathbf{x})\psi_m(\mathbf{x})w(t)$$
$$= \sigma_m(w, a_m)_t \quad (40)$$

where the sum on t now extends over both stimulated responses and blanks, and where $w(t) = 1$ for a stimulated snapshot and $w(t) = 0$ for a blank. Consequently

$$(\phi(\mathbf{x}), \psi_m(\mathbf{x}))_{\mathbf{x}} = \frac{1}{\sigma_m}(w(t), a_m(t))_t \quad (41)$$

and the indicator function may be written as

$$\phi(\mathbf{x}) = \sum_{n=1}^{T} \frac{1}{\sigma_n}(a_n, w)_t \, \psi_n(\mathbf{x}) \quad (42)$$

Since, in general, $(\psi_n, p)_{\mathbf{x}}$ is unknown T cannot be chosen to explicitly minimize the angle between p and ϕ. Instead we have to rely on the implicit assumption that $a_n(t)$ is correlated with the presence or absence of a response to the stimulus manifest in $\psi_n(\mathbf{x})$. As (42) shows, the indicator function method constructs $\phi(\mathbf{x})$ exploiting spatio-temporal correlations between the stimulus and the response. The 'noise' is also spatio-temporally correlated, but in a different manner. The method will fail if there are spatio-temporal correlations in the noise that are correlated with the stimulus.

Equation (42) also makes clear the central role played in this construction by the K-L expansion which associates spatial eigenfunctions with modal coefficients. The indicator function is a superposition of the eigenfunctions and the weight given to each is directly proportional to the correlation of the modal coefficient with the presence of the stimulus.

At first glance the inversion of K, whose eigenvalues decay slowly to zero, would appear to be a problem that would benefit from a gradual regularization, such as Tikhonov regularization [see, for example, Hansen (1992) and Tikhonov and Arsenin (1977)]. However, in practical applications division by zero is not the chief concern. Rather, it is necessary to exclude from the solution the $\psi_n(\mathbf{x})$ which describe correlations due to chance fluctuations (not just pixel noise) between the stimulated and blank datasets. A sharp truncation achieves this more effectively than a smooth filtering – a fact which has been confirmed by numerical experiment.

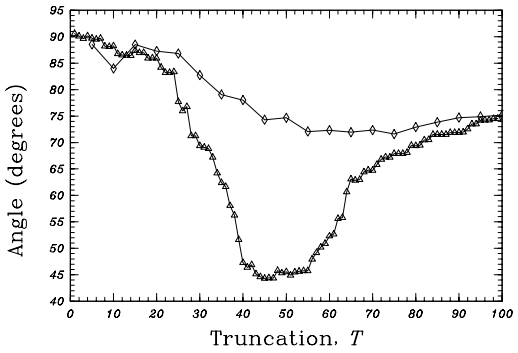

Fig. 8. Angle between the pattern p and the indicator functions of Method I (*triangles*) and Method II (*diamonds*) for contaminated data

5.1 Residuals

The criterion function or residual, as derived from (14) and (18), may be simply expressed as

$$\mathscr{C}(T) = \|w\|^2 - \sum_{n=1}^{T} (a_n, w)_t^2 \quad (43)$$

$$= N - \sum_{n=1}^{T} \left(\sum_{t=1}^{N} a_n(t) \right)^2 \quad (44)$$

where the squared inner sum extends only over stimulated snapshots. Clearly $\mathscr{C}$ decreases as the truncation point, T, increases. If the pattern is represented by only one or two K-L eigenfunctions, $\mathscr{C}(T)$ may be expected to fall sharply at the indices of these eigenfunctions. Figure 9 shows $\mathscr{C}(T)$ for the data described in Sect. 4. The rapid drop which finishes at $T \approx 40$ suggests that there is little advantage to be gained in retaining eigenfunctions beyond this point, and comparison with Fig. 4 shows that $T \approx 40$ is approximately the optimal truncation. $\mathscr{C}(T)$ decreases only very slowly as

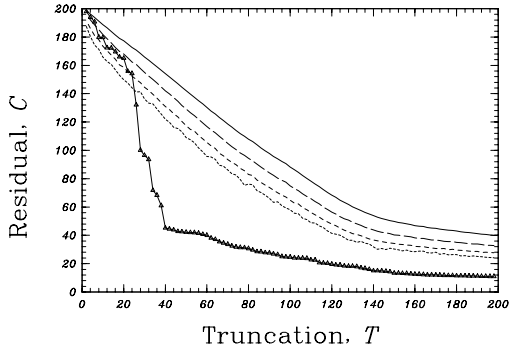

Fig. 9. Variation of residuals, $\mathscr{C}(T)$, with truncation, T. *Symbols* indicate residuals for unshuffled problem. Mean variation of residuals for 10 000 shuffled datasets is given by the *continuous line*. Confidence limits for a single shuffled data set: 0.1 (*long dashes*), 0.01 (*short dashes*) and 0.001 (*dots*)

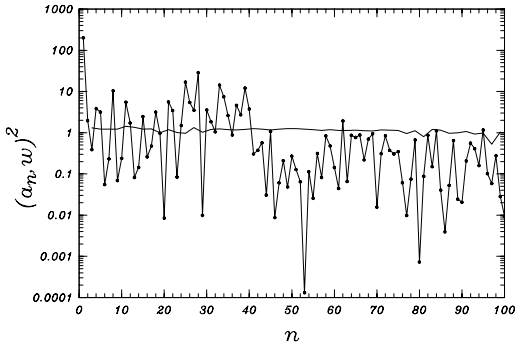

Fig. 10. The quantity $(a_n(t), w(t))$ plotted versus n (*symbols*) and the mean value of $(a_n(t), w(t))$ for shuffled surrogate data

T approaches 150, the point at which the K-L eigenfunctions are dominated by pixel noise. Indicator functions constructed with this many eigenfunctions are very well correlated with the presence or absence of the stimulus, but the spatial structure is overwhelmingly dominated by pixel noise.

The degree to which the indicator function represents a genuine difference between the stimulated and blank datasets can be assessed by comparing $\mathscr{C}(T)$ with residuals derived from shuffled data.

To this end the stimulated and blank datasets are pooled, shuffled and repartitioned into surrogate stimulated and blank datasets, and residuals calculated for the surrogates. Inspection of (43) shows that the residual for the surrogate data is easily calculated by forming the inner products $(a_n, \tilde{w})$, where $\tilde{w} \in \mathbf{R}^{N+M}$ is a vector with ones in N randomly chosen positions and zeroes elsewhere. An ensemble of shuffled residuals permits the estimation of the values below which a given fraction of the residuals fall by chance. Though the shuffled residuals are not (especially for small and large T) normally distributed, these confidence intervals may be estimated by repeated shuffling and explicit generation of the distributions.

Figure 9 also shows the mean shuffled residuals $\tilde{\mathscr{C}}(T)$ and the 0.1, 0.01 and 0.001 confidence contours for single runs of shuffled data. When $T \geq 40$ we can be confident that the indicator function certainly represents a difference between stimulated and blank data. The maximum difference between the mean shuffled residual and $\mathscr{C}(T)$ occurs at T = 38, only slightly smaller than the optimum truncation.

The contribution of the T^{th} eigenfunction to the residual is measured by

$$\mathscr{C}(T-1) - \mathscr{C}(T) = (a_T, w)^2 \tag{45}$$

while its contribution to the indicator function is $(a_T, w) \cdot \psi_T(\mathbf{x})/\sigma_T$. A reasonable choice of truncation point is, therefore, the T for which $(a_T, w)^2$ becomes smaller than $\langle (a_T, \tilde{w})^2 \rangle$, where $\langle \cdot \rangle$ denotes the average over an ensemble of shuffled $\tilde{w}$. As Fig. 10 shows, when $n < 45$ $(a_n, w)^2$ is greater than might be expected from shuffled data, whereas when $n > 45$ the mean contribution from surrogate data is as great as that from the real data, and there is little point in retaining eigenfunctions past $T \approx 45$.

5.2 Reduced problems

The efficacy of a particular indicator may be assessed by projecting it out of all the data (to form 'reduced' datasets) and then seeking another indicator function to discriminate between these reduced datasets. If the first indicator function was successful in capturing the difference between the stimulated and blank datasets, no indicator function (for any truncation) will identify a substantial difference between the reduced stimulated and blank datasets. More precisely, denote by $\phi(\mathbf{x}; T)$ the indicator function formed with truncation T, then the reduced datasets are

$$\begin{aligned} f'(t, \mathbf{x}) &= f(t, \mathbf{x}) - (\phi(\mathbf{x}; T), f(\mathbf{x}, t))_{\mathbf{x}}\ \phi(\mathbf{x}; T)/\|\phi(\mathbf{x}, T)\|^2 \\ \hat{f}'(\hat{t}, \mathbf{x}) &= \hat{f}(\hat{t}; \mathbf{x}) - (\phi(\mathbf{x}, T), \hat{f}(\mathbf{x}, \hat{t}))_{\mathbf{x}}\ \phi(\mathbf{x}; T)/\|\phi(\mathbf{x}, T)\|^2 \end{aligned} \tag{46}$$

The indicator functions for the reduced data at truncations T' may be denoted $\phi'(\mathbf{x}, T')$ and we find the maximum difference between the mean shuffled residuals and the unshuffled residuals for the reduced problem:

$$\rho(T) = \max_{T'} \left\{ \mathscr{C}'(T') - \langle \tilde{\mathscr{C}}'(T') \rangle \right\} \tag{47}$$

In practice it is usually unnecessary to recompute $\langle \tilde{\mathscr{C}}'(T') \rangle$ as it is extremely well approximated by $\langle \tilde{\mathscr{C}}(T)) \rangle$ for the full problem. $\rho(T)$ is shown in Fig. 11 together with residuals for the reduced problems at different T. When T is small $\phi(\mathbf{x}; T)$ performs poorly, while when T is larger than optimum $\phi(\mathbf{x}; T)$ is dominated by chance fluctuations between f and $\hat{f}$, so that in both cases $\phi'(\mathbf{x}, T')$ is able to discriminate between the reduced datasets. Close to the optimum, a truncation which projects $\phi(\mathbf{x}; T)$ out of f and $\hat{f}$ has removed the principal difference between them and $\phi'(\mathbf{x}; T')$ is unable to distinguish between them at any truncation T'. In Fig. 11 this criterion suggests a rather broad range for the truncation. It does, however, span the optimum T and shows when chance fluctuations begin to dominate.

5.3 Subspace support

A further estimate of the best truncation may be obtained by monitoring the degree to which the indicator function

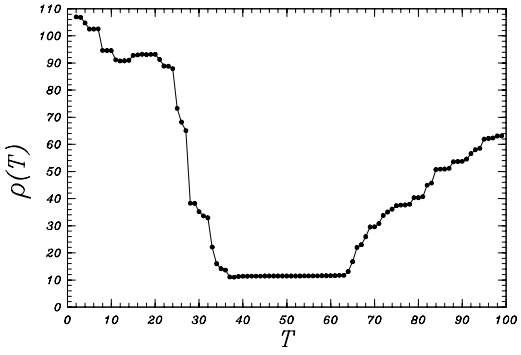

a

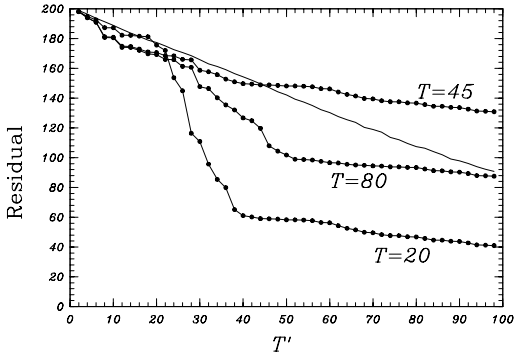

b

Fig. 11. a Maximum difference $\rho(T)$ between shuffled and unshuffled residuals. **b** Residual versus truncation, T', for reduced problems corresponding to T = 20, 45 and 80. The *continuous line* is the mean residual over many shufflings

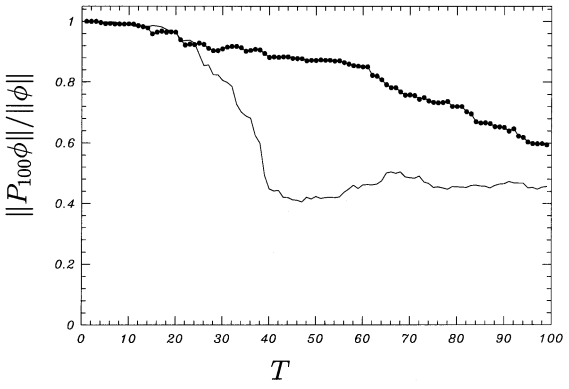

Fig. 12. Proportion of indicator function $\phi(\mathbf{x};T)$ residing in stimulated (*symbols*) and blank (*continuous*) subspaces

resides in each of the stimulated and blank spaces. Define, as above, P_n and $\hat{P}_n$ to be orthogonal projectors onto the principal n K-L eigenfunctions of the stimulated and blank subspaces respectively. Choosing n = 100 (Fig. 3), we then monitor $\|P_n\phi(\mathbf{x};T)\|/\|\phi(\mathbf{x};T)\|$ and $\|\hat{P}_n\phi(\mathbf{x};T)\|/\|\phi(\mathbf{x};T)\|$ as shown in Fig. 12. Ideally $\|P_n\phi(\mathbf{x};T)\|/\|\phi(\mathbf{x};T)\| = 1$ and $\|\hat{P}_n\phi(\mathbf{x};T)\|/\|\phi(\mathbf{x};T)\| = 0$, however, as the figure shows, the proportion lying in the blank space is large until the indicator function can distinguish between the two spaces. The maximum difference occurs for T = 48 which is close to the best truncation.

5.4 Summary

As noted above we are forced to rely on indirect methods to determine the best truncation. Each of the criteria discussed above is helpful in choosing the truncation, but as they are not unequivocal, poor data will encourage the use of more than one of them. In exploratory data analysis confidence intervals estimated from repeated shufflings are frequently helpful in determining whether there is any signal at all.

6 Extensions and variations

In certain situations the response to a particular stimulus can be evoked but with opposite sign by stimulating with a complementary stimulus. This is near the case for ocular dominance columns in the visual cortex. Denoting the data that are expected to contain the response with opposite sign by $\tilde{f}(t,\mathbf{x})$, we may seek an indicator function such that

$$(f(t,\mathbf{x}),\phi(\mathbf{x}))_{\mathbf{x}} = +1, \quad t = 1,\ldots,N$$
$$(\tilde{f}(\tilde{t},\mathbf{x}),\phi(\mathbf{x}))_{\mathbf{x}} = -1, \quad \tilde{t} = 1,\ldots,\tilde{N} \tag{48}$$

Again appealing to a variational principle, we obtain, as might be expected from linearity,

$$\sum_{\mathbf{y}} K(\mathbf{x},\mathbf{y})\phi(\mathbf{y}) = \langle f\rangle - \langle\tilde{f}\rangle \tag{49}$$

The considerations above apply in a straightforward manner to this case. We remark that in this case there is no analogue of the Method II indicator function, because the signal now resides in both f and $\tilde{f}$. If, in addition, unstimulated measurements $\hat{f}(t,\mathbf{x})$ are available, (48) may be augmented with

$$(\hat{f}(\hat{t},\mathbf{x}),\phi(\mathbf{x}))_{\mathbf{x}} = 0, \quad \hat{t} = 1,\ldots,M \tag{50}$$

which leads to an increase in the quality of the resulting indicator function.

It commonly occurs that there are several, say G, groups or classes of stimulation and we wish to find indicator functions $\phi^{(g)}(\mathbf{x})$, $g = 1,\ldots,G$, which best represent the distinguished class. Clearly, if an additional class of blanks is available one may find G indicator functions as described above. Often the classes are, in some sense, on an equal footing, such as the responses to the same stimulus but presented at G different orientations around the circle (for a discussion of symmetries in the visual cortex see Sirovich et al. 1996). In this instance one may use the undistinguished classes as the blank set; that is, we can seek G functions each of which minimizes the criterion function

$$\mathscr{C}^{(g)} = \sum_{f\in S_g}((f,\phi^{(g)})_{\mathbf{x}} - 1)^2 + \sum_{f\notin S_g}(f,\phi^{(g)})^2_{\mathbf{x}} \tag{51}$$

where S_g is the set of responses to the gth stimulus.

It should be noted that there is an approximate linear dependence among the G indicator functions found in this manner. Variation of (51) leads to expressions for $\phi^{(g)}$ in terms of the covariance function of the dataset and the conditional averages $\langle f\rangle_g$ for each of the stimulus classes:

$$\sum_{\mathbf{y}} K_T(\mathbf{x},\mathbf{y})\phi^{(g)}(\mathbf{y}) = N\,\langle f\rangle_g \tag{52}$$

where K_T is the covariance composed of all the classes restricted to the space spanned by the first T K-L eigenfunctions. The sum of the G indicator functions thus satisfies

$$\sum_{\mathbf{y}} K_T(\mathbf{x},\mathbf{y})\sum_{g=1}^{G}\phi^{(g)}(\mathbf{y}) = \sum_{g=1}^{G}\langle f(\mathbf{x})\rangle_g = \langle f(\mathbf{x})\rangle \tag{53}$$

and so

$$\sum_{g=1}^{G}\phi^{(g)}(\mathbf{y}) = \sum_{n=1}^{T}\sigma_n^{-2}\psi_n(\mathbf{y})\sum_{\mathbf{x}}\psi_n(\mathbf{x})\langle f(\mathbf{x})\rangle \tag{54}$$

where the ψ_n are the K-L eigenfunctions for the ensemble comprised of all the snapshots. Writing f in terms of its K-L expansion (2) yields

$$\begin{aligned}\sum_{g=1}^{G}\phi^{(g)}(\mathbf{y}) &= \sum_{n=1}^{T}\sigma_n^{-2}\psi_n(\mathbf{y})\sum_{\mathbf{x}}\psi_n(\mathbf{x}) \qquad (55)\\ &\quad\times\frac{1}{M}\sum_{t=1}^{M}\sum_{k=1}^{T}a_k(t)\sigma_k\psi_k(\mathbf{x})\\ &= \sum_{n=1}^{T}\sigma_n^{-1}\langle a_n\rangle\psi_n(\mathbf{y}) \qquad (56)\end{aligned}$$

Since the $a_n(t)$ are orthonormal the elements of only one of them (usually the first) can be of the same sign. Consequently there is substantial cancellation in the sums $\langle a_n\rangle$ so that $\langle a_n\rangle/\sigma_n \approx 0$. Thus the right-hand side of (56) is approximately, though not exactly, zero and the G indicator functions have an approximate linear dependence. This may be understood by considering that each of the $\phi^{(g)}$ discriminates between the snapshots in it own class and the snapshots in all other classes; consequently, any response common to all classes will not be represented in any $\phi^{(g)}$. In many instances this common mode rejection is distinctly advantageous.

Symmetry considerations may be used to place additional constraints on the G indicator functions (51); the particular application to determining the cortical response to oriented stimuli is discussed in Everson et al. (1997).

Finally, we mention that the quadratic powers appearing in the criterion function (14) may be replaced by a more general power, say q:

$$\mathscr{C}(\phi) = \sum_{t=1}^{N}|(f,\phi)_{\mathbf{x}} - 1|^q + \sum_{\hat{t}=1}^{M}|(\hat{f},\phi)_{\mathbf{x}}|^q \tag{57}$$

Larger values of q emphasize those f which deviate strongly from the mean, while smaller values of q might be expected to provide a degree of robustness to outliers. Numerical schemes to find ϕ in this case are readily implemented; however, we not have found that other values of q provide any substantial improvement over the least squares method which is much more efficiently computed.

7 Summary

We have examined two methods for estimating a small response due to stimulation which is masked by a fluctuating background. Loosely, the first method seeks a indicator function that is most parallel to the stimulated data and most orthogonal to the unstimulated data. The second method, which consists of removing part of the subspace spanned by the unstimulated data, is usually inferior to the first method because the signal is contained, though at smaller amplitude, by the blank subspace as well as the stimulated subspace.

Construction of the Method I indicator function relies on restricting the function to be a superposition of the first T K-L eigenfunctions of the data. We have discussed several criteria by which to choose the optimal T and, as remarked above, poor data will encourage the use of more than one of them.

Use of the Method I indicator function has permitted the detection of cortical structures in poor data for which traditional subtraction procedures were ineffective.

Acknowledgements. Many fruitful discussions with Ehud Kaplan, Estela O'Brien and Darren Orbach are gratefully acknowledged. We are indebted to Jonathan Victor for pointing out the similarities of this approach with generalized indicator functions. Supported in part by grants: NET-EY 11276, NIMH-R0I MH 50166, ONR-N00014-96-1-0492, ONR-N00014-96-1-5005.

References

Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science 273:1868–1871

Blasdel G (1992) Differential imaging of ocular dominance and orientation selectivity in monkey striate cortex. J Neurosci 12:3115–3138

Blasdel G, Salama G (1986) Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. Nature 321:579–585

Duda R, Hart P (1973) Pattern classification and scene analysis. Wiley, New York

Everson R, Kaplan E, Knight B, O'Brien E, Orbach D, Sirovich L (1997) Analysis of optical imaging datasets. (In preparation)

Everson R, Sirovich L, Knight B, Kaplan E, O'Brien E, Orbach D (1997) Optical imaging of the mammalian visual cortex. Bio Bull (in press).

Fisher R (1936) The use of multiple measurments in taxonomic problems. Ann Eugenics 7:179–188

Frostig R, Lieke E, Ts'o D, Grinvald A (1990) Cortical functional architecture and local coupling between neuronal activity and the microcirculation revealed by in vivo high-resolution optical imaging of intrinsic signals. Proc Natl Acad Sci USA 87:6082–6086

Hansen PC (1992) Analysis of discrete ill-posed problems by means of the L-curve. SIAM Rev 34:561–580

Karhunen K (1946) Zur Spektraltheorie Stochastischer Prozesse. Ann Acad Sci Fenn 34:1–7

Loève M (1955) Probability theory. Van Nostrand, Princeton, NJ

Lorenz E (1956) Empirical orthogonal functions and statistical weather prediction. (Technical Report 1) Statistical Forecasting Project, Department of Meteorology, MIT, Cambridge, Mass

O'Brien E, Everson R, Kaplan E (1995) Intrinsic optical imaging of mammalian striate cortex: luminance contrast response, spatial and temporal frequency tuning in co blobs and interblobs. Invest Ophthalmol Visual Sci 36:873

Schmidt E (1907) Zur Theorie der linearen und nichtlinearen Integralgleichungen. I. Entwicklung willkürlicher Funktion nach Systemen Vorgeschriebener. Math Ann 63:433–476

Sirovich L (1987) Turbulence and the dynamics of coherent structures. I: Coherent structures. II: Symmetries and transformations. III: Dynamics and scaling. Q of Appl Math 45:561–590

Sirovich L, Everson R (1992) Analysis and management of large scientific databases. Int J Supercomput Appl 6:50–68

Sirovich L, Everson R, Kaplan E, Knight B, O'Brien E, Orbach D (1996) Modeling the functional organization of the visual cortex. Physica D 96:355–366

Stewart G (1993) On the early history of the singular value decomposition. SIAM Rev 35:551–566

Tikhonov A, Arsenin V (1977) Solutions of ill-posed problems. Wiley, New York

Ts'o D, Frostig R, Lieke E, Grinvald A (1990) Functional organization of primate visual cortex revealed by high resolution optical imaging. Science 249:417–420