Optical Imaging of the Mammalian Visual Cortex

R. Everson, E. Kaplan, B.W. Knight, E.O'Brien, D. Orbach, L. Sirovich

Laboratory of Biophysics, Rockefeller University

&

Laboratory for Applied Mathematics, Mount Sinai Medical School, CUNY,

New York City, NY.

April 29, 1996

Abstract

In vivo optical imaging of the visual cortex (both of intrinsic signals as well as with voltage-sensitive dyes) makes the activity of entire neuronal assemblies accessible for the first time. However, the magnitude of the data collected (>1 Gbyte/minute) as well as the tiny signal-to-noise ratio ($O(10^{-4})$) necessitate the development of optimized analysis algorithms. We have developed extensions of the Karhunen-Loève principal components analysis that effectively filter out unwanted noise and isolate the visually-relevant cortical response; the size of the remaining data set is drastically reduced. The implementation of these techniques on a Parallel Virtual Machine, which allows real-time extraction of cortical functional architecture and dynamics, is described as well. These techniques can be a powerful tool wherever manipulation of a large scientific data base is required.

Keywords: *optical imaging, visual cortex, orthogonal eigenfunctions, data compression, PVM*

Running head: *Optical imaging*

Supported by NIMH R01MH50166-01, ONR N00014-95-1-0281 and N00014-93-1-0279.

Introduction

Light from the external visual world is focused by the lens of the eye onto the retina. Some of the many transformations of visual information take place in the layers of the retina. The first layer of a typical mammalian retina is comprised of $O(10^8)$ photoreceptors, while the final stage of the retina, the ganglion cell layer, sends $O(10^6)$ fibers into the optic nerve. Thus there is an $O(10^2)$ convergence in the visual information collected by the eye. The lateral geniculate nucleus (LGN), which is the next sub-cortical station, also contains $O(10^6)$ neurons. Both the retina and the LGN process information in a circularly symmetric fashion; neurons at the retinal and LGN stages have virtually circular local receptive fields. Information leaves the LGN in (at least) two distinct streams. In brief, the P(Parvo)-stream carries information which is relatively highly resolved in space, is color coded , and is temporally low-pass. The M(Magno)-stream is optimal for coarse grain patterns, is broadband in color and temporally bandpass. As this suggests, the receptive fields of M-cells are significantly larger (4 to 9 times in area) than those of the P-cells, and the former are significantly more sensitive to luminance contrast than are the latter [1].

The next stage in the transformation of visual information occurs at the primary visual cortex, also known as V1 (or area 17). The cortex is a thin piece of tissue (about 1.7 mm thick from dura to white matter) composed of six layers, and V1 occupies an area of about 15 cm^2. Visual information is greatly elaborated in the primary visual cortex, a fact which is underlined by the presence of $O(10^8)$ neurons in V1, rougly 100 times as many as in the LGN. Circular symmetry is replaced by elongated receptive fields for the neurons in the upper layers of the cortex. The familiar modalities of vision such as color, texture, orientation, motion and so forth, are represented in the cortex, and this must be so for every visible position in visual space. It is found that neurons subserving a particular modality are organized in *columns* perpendicular to the surface of the cortex. In addition, for many modalities, cells with similar eye, orientation or motion preference tend to be clumped in *patches.*

This picture of the architecture and functionality of the visual cortex is based on over 25 years of detailed experiments using single unit recording and staining methods [2, 3].

Optical Imaging

A new chapter in neurophysiological research began with the introduction of optical imaging methods. Cohen and his colleagues [4] introduced the use of voltage sensitive dyes to monitor neuronal activity. (For a review see Lieke *et al* [5].) Such dyes, which respond instantly to changes in membrane voltage, permit the polling of vast populations of neurons, in contrast to the single unit recordings from which our present knowledge of cortical dynamics is largely derived. A related form of imaging, known as intrinsic signal imaging, makes use of changes in the reflectance of cortical tissue due to metabolic activity induced by neuronal activity, and without the use of dyes. The discovery of the intrinsic signal by Grinvald *et al* [6], allows measurements of activity, free of dye related problems such as phototoxicity. In both instances responses to visual stimuli are extremely small. For example, the intrinsic signal measures cortical reflectance which includes contributions from normal (unstimulated) background activity, such as cardiac and respiratory activity, as well as noise. The sought-after signal is $O(10^{-3})$ or less of the total signal. Figure 1 shows four cortical images collected while a macaque viewed four *different* stimuli: differences in the cortical response are imperceptible to the naked eye.

Optical imaging offers the possibility of going well beyond the local activity of individual and neighboring neurons. We can now attempt to measure assemblies of neurons engaged in correlated activity and reveal the dynamics and underlying network of connections at work in the cortex.

Subtraction Methods

A standard procedure for extracting the extremely small optical signal is referred to as the method of subtraction. For purposes of background and later comparison we briefly discuss this procedure.

In general, an image will be represented by the gray level, $\phi(\mathbf{x})$, at pixel location $\mathbf{x}$. More particularly, we denote the optical response to visual stimulation of the left eye by $\phi_\ell = \phi(\ell;\mathbf{x})$. (The particular stimulus is not of concern at this moment.) We can write this as

$$\phi_\ell = \phi_B(\mathbf{x}) + \varphi(\ell;\mathbf{x}), \tag{1}$$

where ϕ_B represents the background (vegetative) signal. Similarly, the right eye response can be written as

$$\phi_r = \phi_B(\mathbf{x}) + \varphi(r;\mathbf{x}). \tag{2}$$

In each instance φ represents the level of activity above the background. The background is eliminated by subtraction:

$$\Delta\phi = \phi_r - \phi_\ell = \varphi(r;\mathbf{x}) - \varphi(\ell;\mathbf{x}). \tag{3}$$

To minimize effects of noise and drift one averages over an ensemble of such differences,

$$R_{oc} =< \Delta\phi >= \frac{1}{N}\left(\sum \phi_r - \sum \phi_\ell\right), \tag{4}$$

where N is the number of pairs of measurements. R_{oc} is referred to as the ocular dominance map, and an example is shown in Figure 2. Here, the light gray regions are active when the left eye is stimulated, corresponding to a decrease in reflectance, while the dark gray regions are active during right eye stimulation. Ocular dominance maps obtained in this way show good qualitative agreement with similar maps obtained by staining techniques (for a review see LeVay & Nelson [3]).

A similar program can be carried out for orientation maps. Drifting bars, say of orientation θ, are a standard visual stimulus. (For purposes of exposition we do not discuss here the separation of direction and orientation.) Cortical cells are known to be narrowly tuned to orientation [7]. Consequently, cells which fire for an orientation θ_0, are usually relatively silent for $\theta = \theta_0 \pm \pi/2$. Hence, if we represent the response image to bars of orientation θ by

$$\phi(\theta;\mathbf{x}) = \phi_B(\mathbf{x}) + \varphi(\theta;\mathbf{x}) \tag{5}$$

we can eliminate the background by considering

$$\Delta_\theta \;=\; \{\phi(\theta;\mathbf{x}) + \phi(\theta+\pi;\mathbf{x})\} \;-\; \left\{\phi\left(\theta+\frac{\pi}{2};\mathbf{x}\right) + \phi\left(\theta-\frac{\pi}{2};\mathbf{x}\right)\right\} \tag{6}$$

As before we can form the ensemble average

$$R_\theta =< \Delta_\theta > . \tag{7}$$

An example of such a response is exhibited in Figure 3. The dark regions corresponds to maximum response (a decrease in cortical reflectance) at θ (or $\theta+\pi$) and the dark regions to maximum response at $\theta \pm \pi/2$. Through a systematic investigation of the response patches one can generate an orientation preference map [8].

Subtraction methods are, however, susceptible to fluctuations in the background reflectance, ϕ_B, resulting from changes in the animal's physiology. There may be imperfect cancellation of the background, making it difficult to distinguish the response elicited by stimulation. Below we outline a new technique for extracting the response in an optimal manner.

Eigenfunction Methods

Optical imaging of the cortex generates extremely large databases. Images containing as many as $O(10^6)$ pixels can be produced at the rate of $O(10^2)$ per second. To accommodate the low signal to noise ratio a dynamic range of 2^{12} gray levels is required. A minute of recording can therefore produce over five gigabytes of data. The combination of large data sets and low signal to noise underline the need for innovative approaches for the analysis and interpretation of the results.

From a number of perspectives the Karhunen-Loève (KL) procedure offers an ideal framework for examining such data. (See Sirovich & Everson [9] for an earlier presentation of such a cortical study.) It creates a coordinate system which is intrinsic to the data, and is optimally compact. For completeness we briefly review KL theory.

In general, an optical image can be represented by

$$\phi = \phi(t, \mathbf{x}) \tag{8}$$

where ϕ is the gray level at pixel location $\mathbf{x}$ at time t. For purposes of presentation it is convenient to regard both t and $\mathbf{x}$ as continuous. We seek orthonormal functions $\{\psi_n(\mathbf{x})\}$ in space and $\{a_n(t)\}$ in time together with constants $\{\mu_n\}$ such that

$$\phi(t, \mathbf{x}) = \sum_n a_n(t)\mu_n\psi_n(\mathbf{x}) \tag{9}$$

where

$$\int \psi_n(\mathbf{x})\psi_m(\mathbf{x})d\mathbf{x} \;=\; (\psi_n, \psi_m)_{\mathbf{x}} \;=\; \delta_{nm} \tag{10}$$

and

$$(a_n, a_m)_t = \int a_n(t) a_m(t) dt = \delta_{nm} \tag{11}$$

From (9) it follows that

$$\mu_n a_n = (\psi_n, \phi)_{\mathbf{x}} \tag{12}$$

and

$$\mu_n \psi_n = (a_n, \phi)_t. \tag{13}$$

It follows from these relations that ψ_n is an eigenfunction of the two-point correlation

$$K(\mathbf{x}, \mathbf{y}) = (\phi(t, \mathbf{x}),\ \phi(t, \mathbf{y}))_t, \tag{14}$$

that is

$$\int K(\mathbf{x}, \mathbf{y}) \psi_n(\mathbf{y}) d\mathbf{y} = \mu_n^2 \psi_n(\mathbf{x}). \tag{15}$$

Once ψ_n and μ_n are determined from (15) then $a_n(t)$ is determined from (12). The dual problem of first determining the $\{a_n(t)\}$ is often simpler and is known as the *snapshot method* [10]. Since $K(\mathbf{x}, \mathbf{y})$ is symmetric the eigenfunction problem (15) can always be solved and this, in brief, demonstrates the existence of the expansion (9), which, in fact, is the continuous version of the Singular Value Decomposition [11]. The technique is generally known as the Karhunen-Loève procedure [12, 13]. The continuous form goes back to Schmidt [14], while the discrete version may be identified with Principal Components Analysis [15] and is closely related to factor analysis [16]. The spatial eigenfunctions, $\psi_n(\mathbf{x})$, are known variously as *empirical eigenfunctions*, *principal components* or *empirical orthogonal functions* [17].

If the eigenfunctions are arranged in descending order of the variances μ_n^2, then it follows that in approximating the dataset with an N term truncation of (9), for the choice of our eigenfunctions (15) the mean squared error

$$E_N = \int d\mathbf{x}dt \; \{\phi(t,\mathbf{x}) \; - \; \sum_{n=1}^{N} a_n(t)\mu_n\psi_n(\mathbf{x})\}^2 \tag{16}$$

is minimum for all values of N and over all such eigenfunction expansions. By representing the dataset in terms of the first N empirical eigenfunctions one captures, on average, a fraction $\sum_{n=1}^{N}\mu_n^2/\sum_{n=1}^{\infty}\mu_n^2$ of the variance, the maximum that can be captured using any N orthogonal basis functions.

In Figure 4 we show the first eight eigenfunctions from an experiment on the visual cortex of a macaque monkey. The first of these (having the largest variance) closely resembles the snapshots of Figure 1, and can be associated with background activity. Comparison with vascular images shows that eigenfunctions 3 and 4 are due to vascularity. By contrast, the second eigenfunction is virtually identical to the ocularity map shown in Figure 2, which was obtained for this data set by the subtraction method. This observation lends strong support for the methodologies we have adopted. For the moment, we forego discussion of the remaining eigenfunctions which are associated with orientation and other aspects of the stimulus.

Noise and Background

In Figure 5 we exhibit in log-log form the variance or eigenvalue spectrum, μ_n^2, for the data set under discussion. The first eigenvalue, μ_1^2, is associated with the first eigenfunction which represents the gross reflectance of the cortex: its size relative to the other eigenvalues underlines the tiny size of the intrinsic signal. Apart from μ_1^2, the spectrum is comprised of two different regions.

The final leg, consisting of many small eigenvalues, can be associated with pixel noise; the corresponding eigenfunctions are fluctuations, uncorrelated at all scales. As we have already observed, eigenfunctions 1, 3, and 4 describe background physiological (but probably not neural) activity.

In Figure 6a we show an image at some time t with the background removed, *viz.,*

$$\phi(t, \mathbf{x}) \; - \sum_{n=1,3,4} (\psi_n, \phi)_{\mathbf{x}} \, \psi_n(\mathbf{x}). \tag{17}$$

The animal was viewing a stimulus of drifting random dots with its left eye, and the responding columns are apparent. In Figure 6b we show

$$\tilde{\phi} = \phi(t, \mathbf{x}) - \sum_{\substack{n=1,3,4 \\ n>11}} (\psi_n, \phi)_{\mathbf{x}} \, \psi_n(\mathbf{x}).,\tag{18}$$

i.e., the same image but with the pixel noise also removed. When $\tilde{\phi}$ is viewed in a time sequence, we obtain a dynamical picture of cortical activity in response to the stimuli presented to the animal.

The relative ease with which we were able to analyze the above data set is not typical. Isolation of a modality such as ocular dominance by means of a single eigenfunction, while appealing and desirable, depends greatly on the quality of the data. As already mentioned, the response to a stimulus can be $O(10^{-4})$ of the measured signal. In addition to a daunting signal to noise ratio we must also contend with drift in both the animal's physiology and in the apparatus.

With these points in mind we observe that KL does not yield us a road map for data analysis. Rather it provides us with an unsurpassed view of the terrain. With the use of KL we are generally able to remove both pixel noise and background activity. We skip over some of the technical issues that enter into this process (see Everson *et al* [18]). The images that we discuss below should be regarded as having the background and noise removed. In a manner of speaking the KL procedure furnishes us with a band-pass filter.

Indicator Functions

In order to extract the desired signal we have introduced a form of conditional averaging which leads to the concept of an indicator function [18]. We denote by $\phi_k \in P$ the responses measured in response to a particular stimulus P. The indicator function, $f_P(\mathbf{x})$, is determined by extremizing the criterion function,

$$C = \sum_{\phi_k \in P} ((\phi_k(\mathbf{x}), f_P(\mathbf{x}))_{\mathbf{x}} - 1)^2 + \sum_{\phi_k \notin P} (\phi_k(t, \mathbf{x}),\ f_P(\mathbf{x}))^2_{\mathbf{x}}. \tag{19}$$

Thus, ideally, f_P lies parallel to all responses to the distinguished stimulus and perpendicular to the remaining responses.

In Figure 7 we show ocular dominance columns obtained by subtraction and those obtained by forming

$$f_0 = f_R - f_L, \tag{20}$$

where R and L refer to right and left. With the same goal, alternatively we may let f_0 extremize

$$\sum_{\phi_k \in R} (\phi_k, f_0)_{\mathbf{x}} - 1)^2 + \sum_{\phi_k \in L} ((\phi_k, f_0)^2_{\mathbf{x}} + 1)^2. \tag{21}$$

We point out that the vascularity which mars the subtraction picture is much reduced in the indicator function picture. Also, the indicator function shows zero response along the bottom left edge, which is area V2 and does not contain ocular dominance columns. As the figure illustrates, the procedures work quite well, even in situations when the subtraction methods fail or produce only a faint signal.

As another illustration of the indicator function idea we consider responses to drifting, oriented bars. In the actual experiment 16 directions

$$\theta_k = 2\pi k/16, \quad k = 1, \ldots 16 \tag{22}$$

were used. We therefore obtain sixteen indicator functions, $f_k(\mathbf{x})$, from an expression similar to (19). For simplicity $\phi(\theta_j, \mathbf{x})$ can be regarded as the mean response to bars at orientation θ_j. A typical indicator function, derived from cat data, is shown in Figure 8.

An invariance argument, which is given elsewhere [19], indicates that the characteristic function for orientation is given by the following Fourier transform of the indicator functions:

$$F_2(\mathbf{x}) \ = \ \sum_k f_k(\mathbf{x}) \ e^{-2i\theta_k}. \tag{23}$$

Thus the real and imaginary parts of $F_2(\mathbf{x})$ represent characteristic functions, which in theory have equal variances. Alternatively, we can consider the amplitude and phase of $F_2(x)$, which are shown in Figure 9. As a little thought reveals, the phase plot gives the orientation to which neurons at a $\mathbf{x}$ are tuned. The amplitude portion of the figure gives us the magnitude of the orientation response at $\mathbf{x}$. Notable among the features of the orientation map are the singular points which mark the intersection of many separate orientation regions, and at which the amplitude must fall to zero. The structures surrounding a singular point are referred to as *pinwheels* in the literature [8].

Computational Aspects

The combination of large scale data acquisition and relatively strenuous analysis methods creates some unique computational demands. In assessing the results of imaging methods it is important to be able to make comparison with electrode recordings – still considered to be the *gold standard*. Therefore, the computational demands of experiment and analysis are considerably heightened by the desire to achieve near real time results. In this way we are able to record electrically from cortical regions that are functionally identified by imaging.

Electrical recording of three types can be correlated with the optical images: single unit recording, recordings from small neuronal clusters (2-4 units), and gross potentials or local field potentials. The time series of electrical recordings of each type can be cross-correlated with the time courses of the first few significant eigenfunctions to determine which of them reflects the activity of the neuron(s) from which we recorded.

We use a Parallel Virtual Machine (PVM) [20, 21] to perform real-time Karhunen-Loève decompositions of images as they are recorded. The procedure consists of two phases: a) the empirical eigenfunctions, $\psi_n(\mathbf{x})$, for the particular experiment are determined from an initial experimental sequence; b) the modal coefficients, $\mu_n a_n(t)$, are found for each snapshot by projecting the picture, $\phi(\mathbf{x}, t)$, onto each of the empirical eigenfunctions in turn.

The key to this procedure is the compression provided by the KL decomposition. As was outlined above, the empirical eigenfunctions, $\psi_n(\mathbf{x})$, for large n describe only pixel noise and random fluctuations in the dataset. In our experiments the number of relevant eigenfunctions, with variance above the pixel noise, is usually no more than 100. Therefore we need only find the 100 principal eigenfunctions (in phase a) and form the first 100 modal coefficients (in phase b) to capture all the relevant activity.

This degree of compression, uniquely afforded by the KL decomposition (recall equation (16)), vastly reduces the burden of computation and storage. Elementary estimates show that over a gigabyte of data can be generated in the course of ten minutes of optical recording. Even if it were necessary to monitor 10^3 eigenfunctions, which would require a gigabyte of memory, then the storage of the associated modal coefficients for a ten minute period would need only an additional megabyte of storage. (Longer experimental runs lead to memory requirements which rise linearly, i.e., 10^2 minutes would need 10 megabytes.) While technically this compression is lossy, the nominal thousand modes lie deep in the pixel noise, so that it is noise, and not real information, which is lost.

The PVM in our laboratory consists of ten high-performance CPUs residing in several heterogenous workstations, and linked by a network which allows rapid data transfer and multicasting of data among them. The parallel virtual machine configuration is suited for computing problems which possess an inherent coarse-grained parallelism. Both phases of the procedure we describe have a natural parallelism.

In the initial phase, images passed to the PVM from the camera are multicast to the component machines which assemble the correlation matrix, $C(t,s) = (\phi(t,\mathbf{x}), \phi(s,\mathbf{x}))_{\mathbf{x}}$ and diagonalize it to find the empirical eigenfunctions. This stage is naturally coarse-grained in that each component CPU of the PVM is responsible for computing several columns of the correlation matrix.

The second phase of the procedure involves projecting each snapshot onto the empirical eigenfunctions. Each snapshot recorded by the camera is multicast to a member of the PVM, and that member finds the projections for a group of eigenfunctions. This stage too is naturally parallel, since each projection operation is orthogonal to the others, allowing a component machine to hold, in core memory, the eigenfunctions for which it is responsible, and to project incoming images onto those eigenfunctions without communicating with its peer machines.

The computational bottleneck is currently the transmission of images from the camera to the various machines in the PVM. We anticipate that the installation of a fast FDDI or ATM network, together with (lossless) compression algorithms adapted for the cortical images, will allow us to cope with the very high data rates possible with voltage-sensitive dye imaging.

Summary

Optical imagining of the cortex produces a number of problems of interest to biological imaging and, more generally, to the study of scientific data bases. We encounter data sets which easily exceed a gigabyte. The obvious problem of storage requires sensible compression procedures. Moreover these should be of a nature conducive to analysis. In addition we have to deal with an extremely low signal to noise ratio.

The suite of extensions of the Karhunen-Loève procedure which we have developed meet the above problems. Through the use of the empirical eigenfunctions we are able to filter out noisy and unwanted parts of the data. We are then left with a greatly compressed data set which contains the essential information contained in the original data, and the data are rendered in a form that is conveniently navigated and browsed. The methods and algorithms, while computationally intensive, are nevertheless within reach of available equipment.

References

[1] E. Kaplan and R.M. Shapley. The primate retina contains two types of ganglion cells, with high and low contrast sensitivity. *Proc. Natl. Acad. Sci. USA*, 83:2755–2757, 1986.

[2] D.H. Hubel and T.N. Wiesel. Functional architecture of macaque monkey visual cortex. *Proc. Roy. Soc. Lond. B*, 198(1-59), 1977.

[3] S. LeVay and S.B. Nelson. Columnar organization of the visual cortex. In A.G. Leventhal, editor, *Vision and Visual Dysfunction (The Neural Basis of Visual Function)*, pages 266–315. Macmillan Press, 1991.

[4] L.B. Cohen, H.V. Davila, D. Landowne, A.S. Wagonner, and C.H. Wang. Changes in axon fluorescence during activity: molecular probes of membrane potential. *J. Membr. Biol.*, 19:1–36, 1974.

[5] E.E Lieke, R.D. Frostig, A. Arieli, D.Y. Ts'o, R. Hildesheim, and A. Grinvald. Optical imaging of cortical activity: Real-time imaging using extrinsic dye-signals and high resolution imaging based on slow intrinsic-signals. *Annu. Rev. Physiol.*, 51:543–559, 1989.

[6] A. Grinvald, E. Lieke, R.D. Frostig, C.D. Gilbert, and T.N. Wiesel. Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature*, 324:361–364, 1986.

[7] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond)*, 160:106–154, 1962.

[8] T. Bonhoeffer and A. Grinvald. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature*, 353:429–431, 1991.

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

[10] L. Sirovich. Turbulence and the dynamics of coherent structures, pt i: Coherent structures, pt ii: Symmetries and transformations, pt. iii: Dynamics and scaling. *Quarterly of Applied Mathematics*, XLV(3):561–590, 1987.

[11] G.H. Golub and C.F. Van Loan. *Matrix Computations.* North Oxford Academic, Oxford, 1983.

[12] K. Karhunen. Zur Spektraltheorie Stochasticher. *Prozesse Ann. Acad. Sci. Fennicae*, 37, 1946.

[13] M.M. Loève. *Probability Theory.* Van Nostrand, Princeton, N.J., 1955.

[14] E. Schmidt. Zur Theorie Der Linearen Und Nichtlinearen Integralgleichungen. I Teil: Entwicklung WillküLicher Funktion Nach Systemen Vorgeschriebener. *Mathenatische Annalen*, 63:433–476, 1907.

[15] H. Hotelling. Analysis of complex statistical variables in principal components. *J. Educ Psy.*, 24:417–441, 498–520, 1933.

[16] H. Harman. *Modern Factor Analysis.* Univ. Chicago Press, Chicago, 1960.

[17] E.N. Lorenz. Empirical orthogonal functions and statistical weather prediction. Technical Report 1, Statistical Forecasting Project, Department of Metrology, MIT, Cambridge, 1956.

[18] R.M. Everson, L. Sirovich, B.W. Knight, E. Kaplan, E. O'Brien, and D. Orbach. Indicator functions for isolating functional signals from optical imaging of the mammalian visual cortex. 1995. *(in preparation)*.

[19] L. Sirovich, R.M. Everson, E. Kaplan, B.W. Knight, E. O'Brien, and D. Orbach. Indicator functions for isolating functional signals from optical imaging of the mammalian visual cortex. *Physica D*, 1995. *(submitted)*.

[20] V. S. Sunderam. PVM: a framework for parallel distributed computing. *Concurrency: Practice & Experience*, 2(4):315–339, December 1990.

[21] G. A. Geist, A. Begeulin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam. *PVM: Parallel Virtual Machine – A users' guide and tutorial for networked parallel computing*. Scientific and Engineering Computation. MIT Press, 1994.

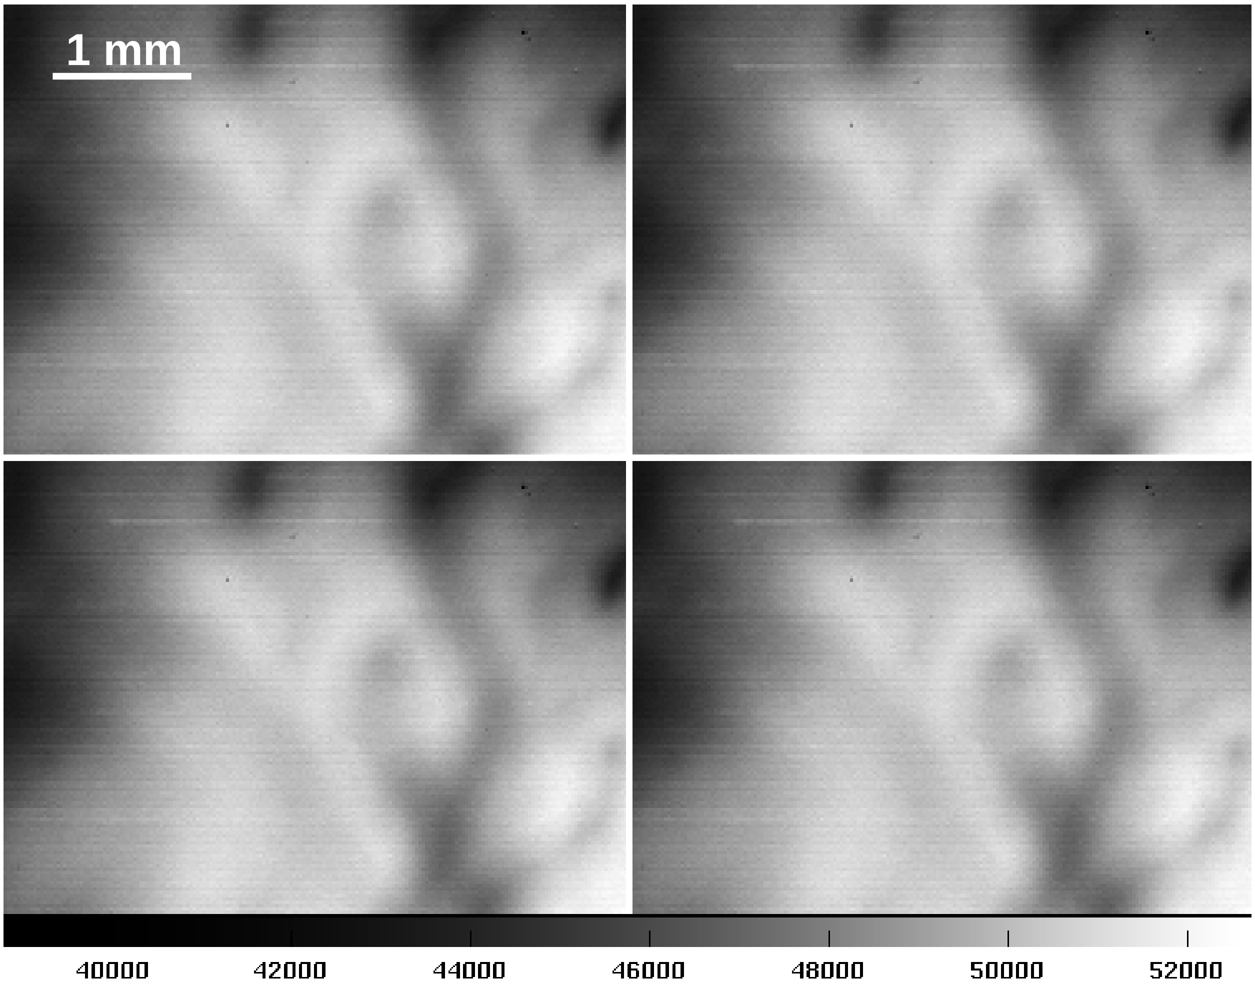

Figure 1: Four intrinsic signal images of area V1 in a macaque monkey. The gray scale indicates the reflectance of the cortex, and although the monkey was viewing a different stimulus as each image was recorded, stimulus-induced reflectance changes are too small to be visible to the naked eye. Part of the vascular tree is evident on the right hand side of the images.

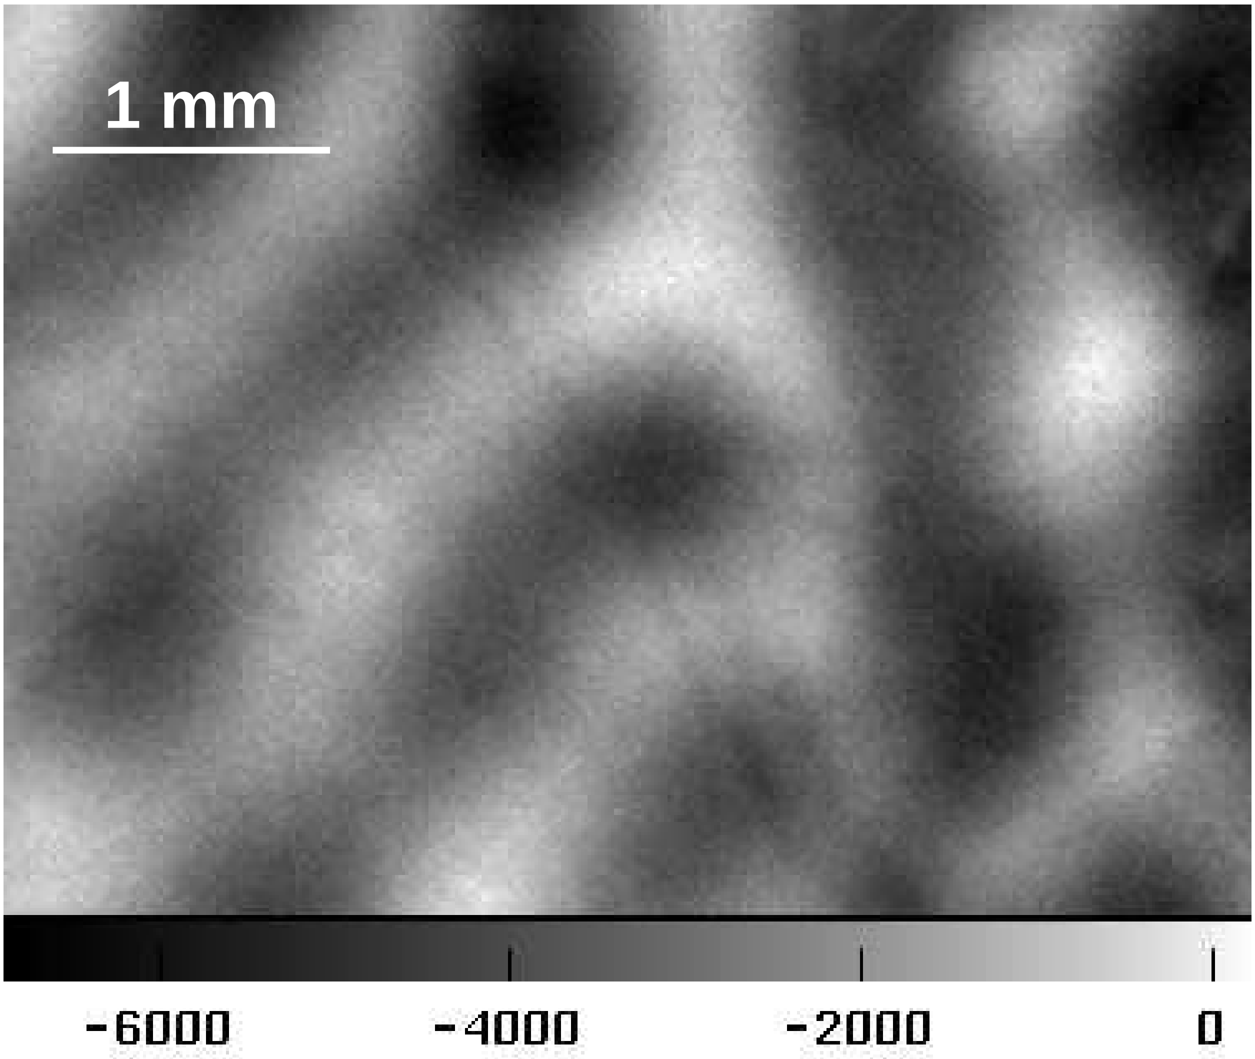

Figure 2: Ocular dominance columns in the macaque obtained by subtracting of images obtained in response to stimulating the right eye from those obtained in response to left eye stimulation.

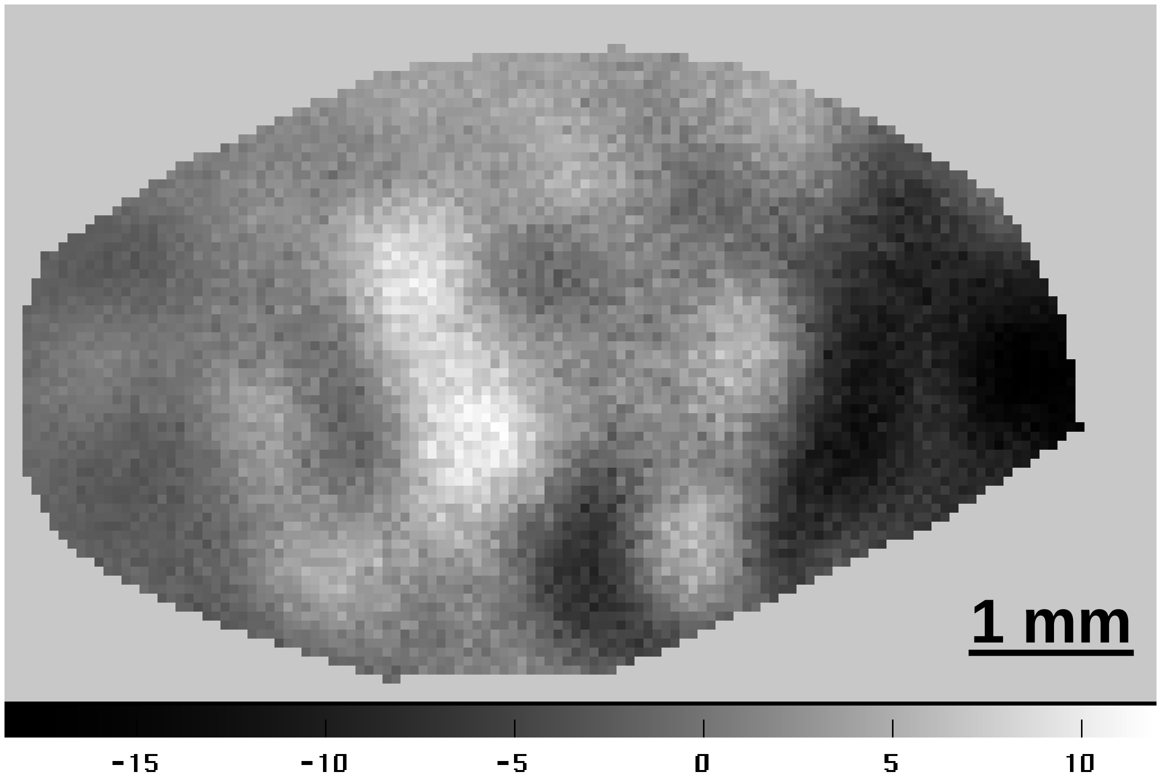

Figure 3: Orientation response in area 18 of a cat found by subtraction of responses obtained to bars drifting in orthogonal directions.

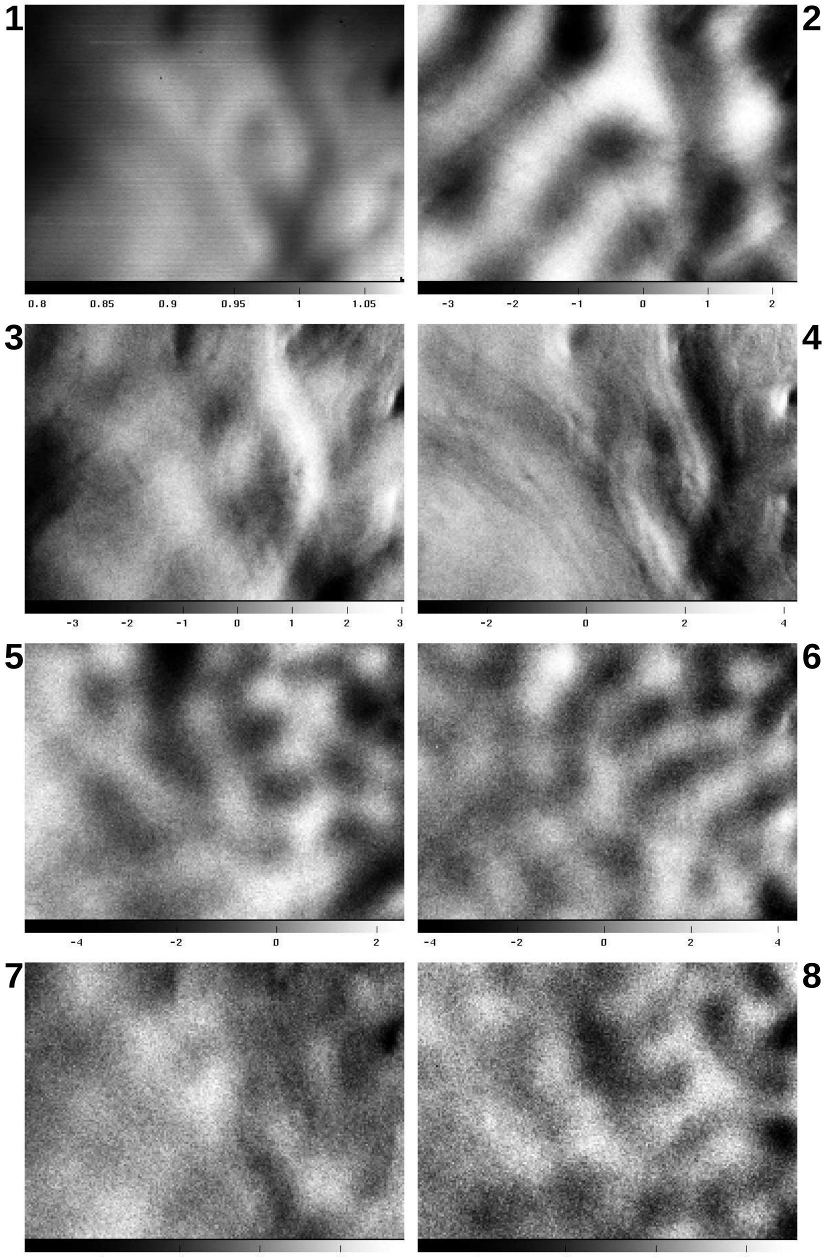

Figure 4: The first eight eigenfunctions, $\psi_n(\mathbf{x})$ from a KL decomposition of intrinsic signal images from a macaque cortex.

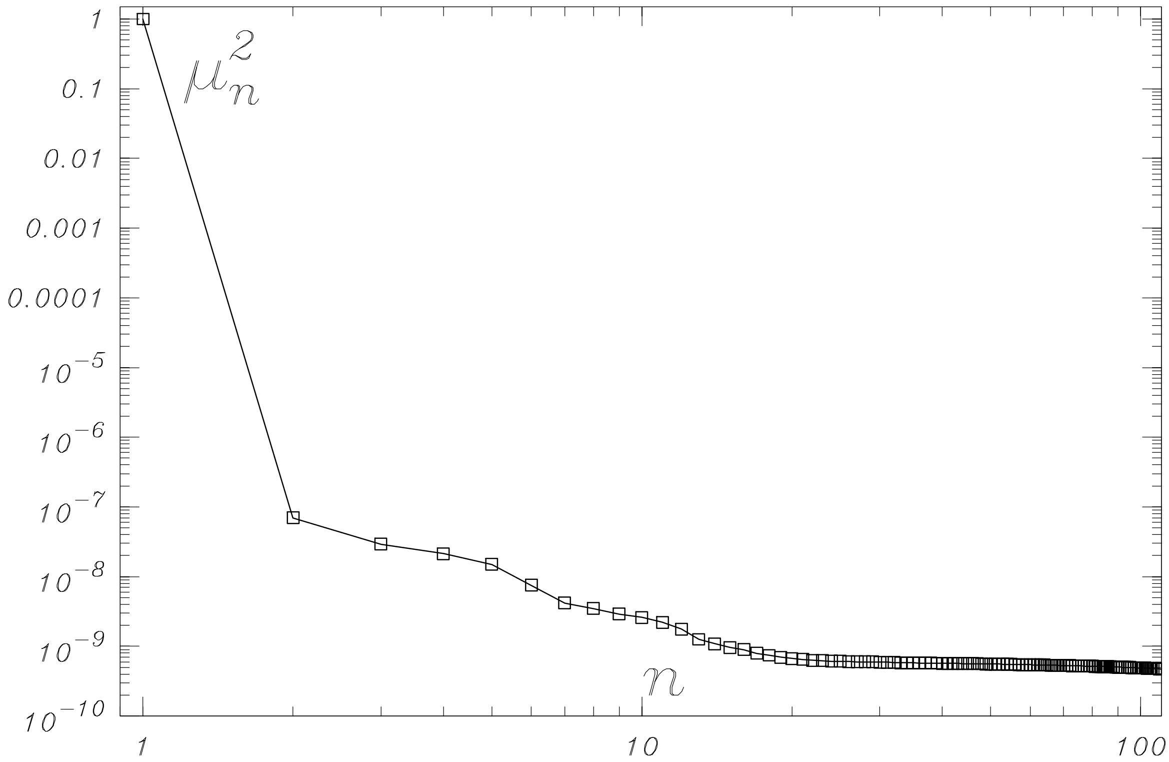

Figure 5: The eigenvalue spectrum associated with the eigenfunctions shown in Figure 4. Each eigenvalue μ_n^2 indicates the mean contribution of the eigenfunctions $\psi_n(\mathbf{x})$ and $a_n(t)$ to representing the dataset. The eigenvalues are normalized so that $\sum_n \mu_n^2 = 1$.

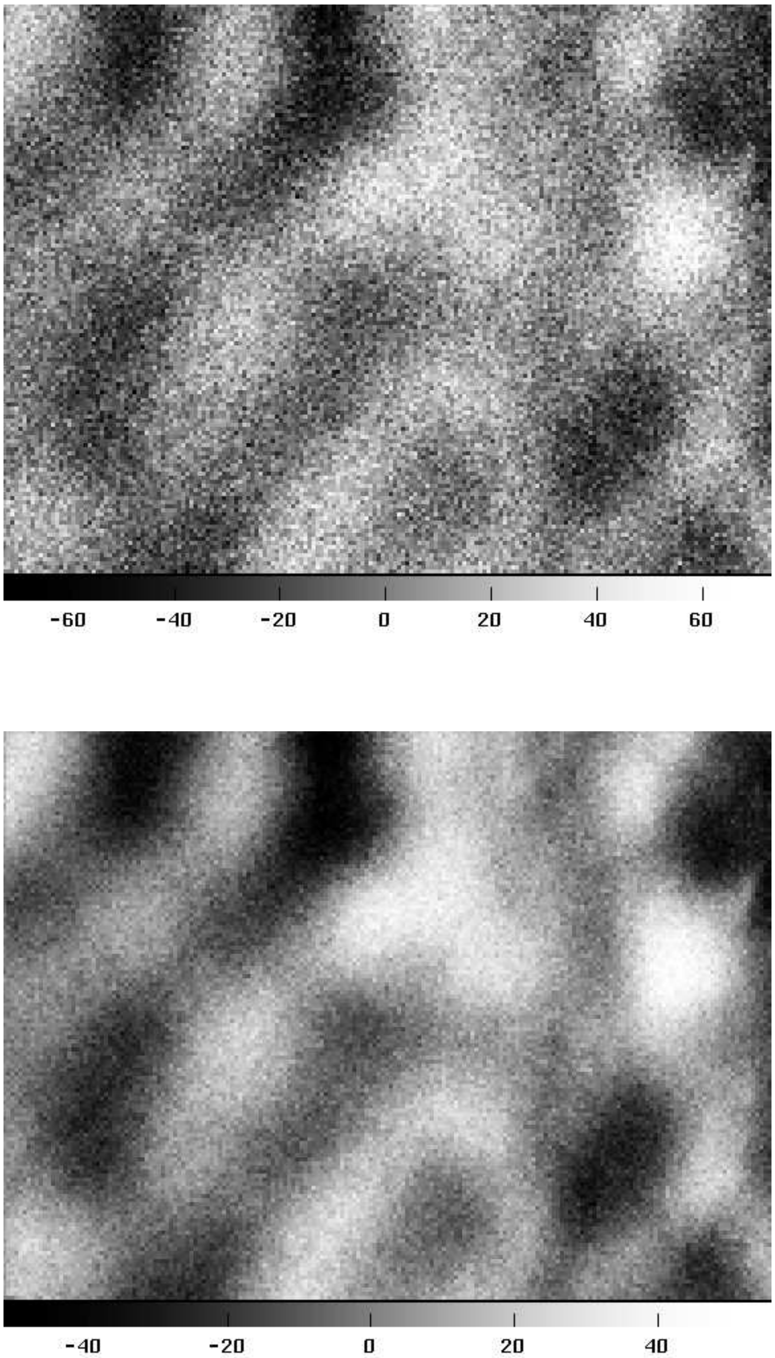

Figure 6: A single snapshot with contributions from eigenfunctions describing noise and vascularity removed. (a) Eigenfunctions 1, 3 & 4, describing the gross reflectance and vascularity subtracted; (b) eigenfunctions ψ_n, $n > 11$, describing pixel noise subtracted as well as those in (a).

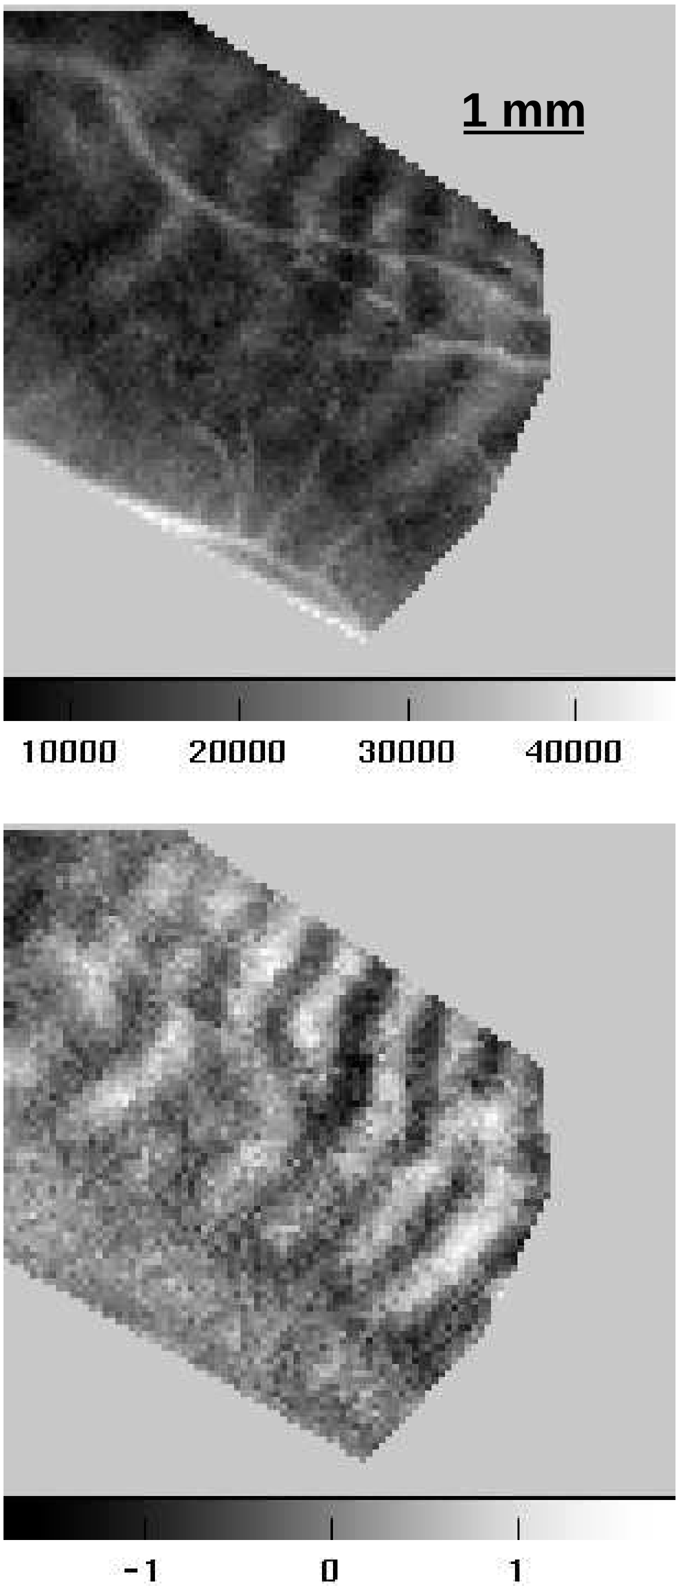

Figure 7: Comparison of ocular dominance columns in a macaque found by subtraction methods and the indicator function technique.

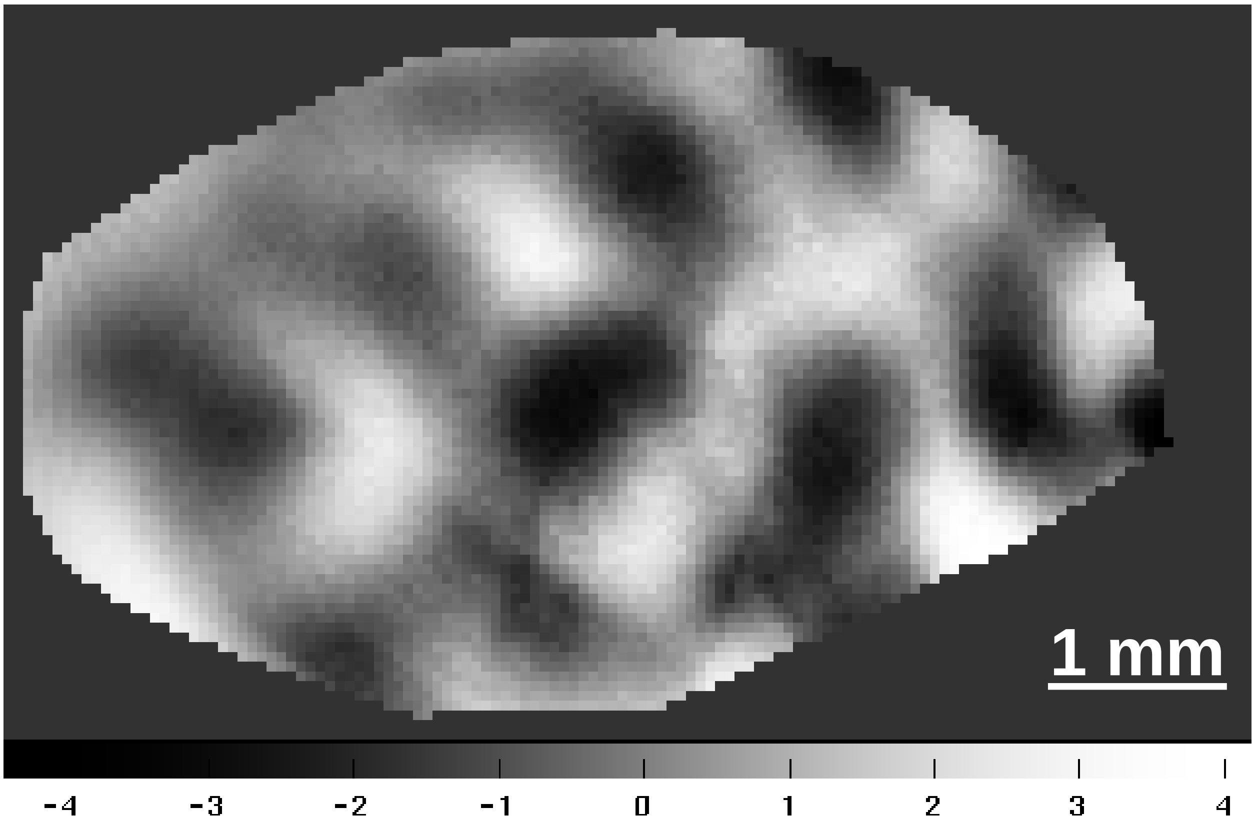

Figure 8: Orientation indicator function from area 18 of a cat.

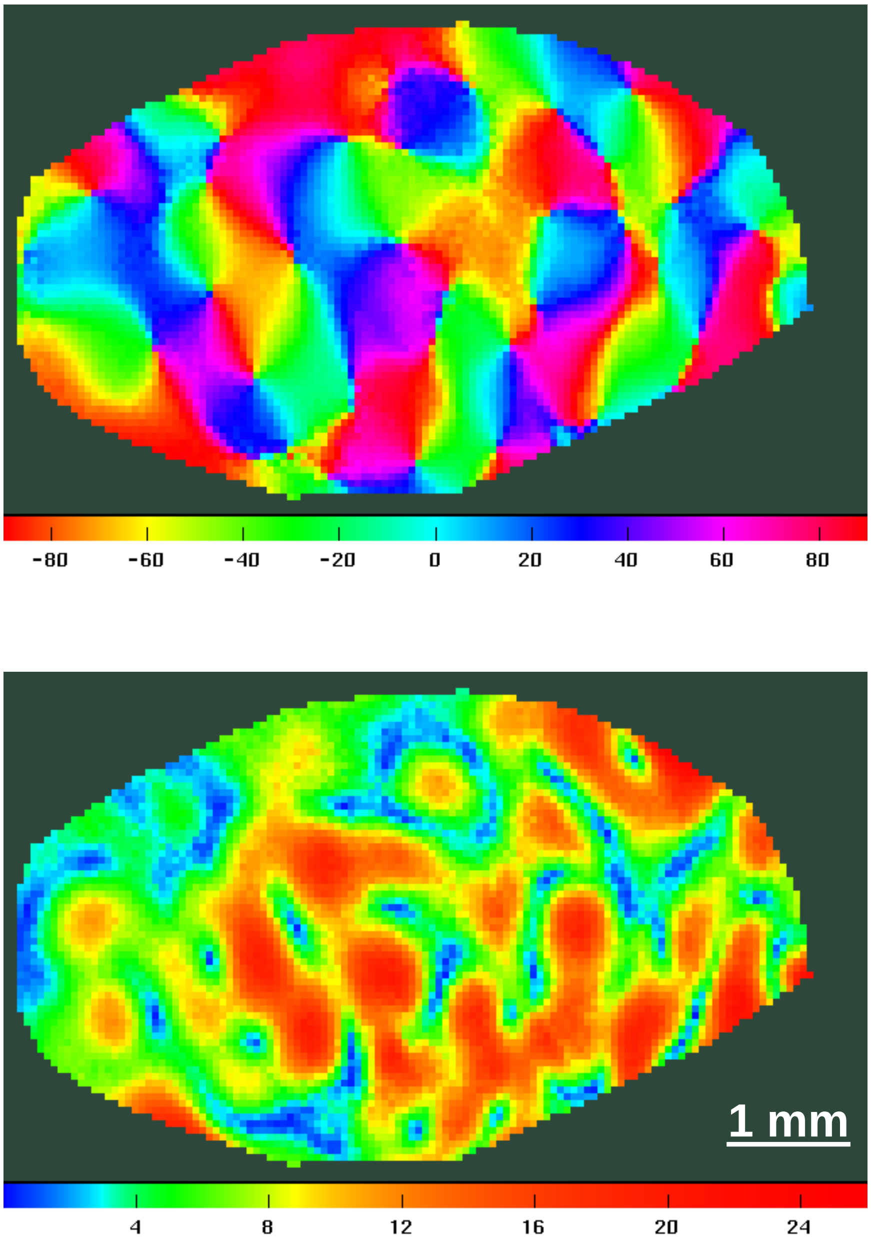

Figure 9: Fourier analysis of orientation indicator functions. The top picture shows the orientation to which neurons are tuned, while the bottom picture gives the magnitude of their response.