

# Analyzing receptive fields, classification images and functional images: challenges with opportunities for synergy

Jonathan D Victor

**In neurophysiology, psychophysics, optical imaging and functional imaging studies, the investigator seeks a relationship between a high-dimensional variable, such as an image, and a categorical variable, such as the presence or absence of a spike or a behavior. The usual analysis strategy is fundamentally identical across these contexts—it amounts to calculating the average value of the high-dimensional variable for each value of the categorical variable and comparing these results by subtraction. Though intuitive and straightforward, this procedure may be inaccurate or inefficient and may overlook important details. Sophisticated approaches have been developed within these several experimental contexts, but they are rarely applied beyond the context in which they were developed. Recognition of the relationships among these contexts has the potential to accelerate improvements in analytic methods and to increase the amount of information that can be gleaned from experiments.**

Many systems neuroscience experiments are based around a common basic design—identifying an association between a high-dimensional variable, such as a complex stimulus, and a variable that can be easily categorized, such as the presence or absence of neural spiking. For example, in receptive field analysis, the investigator presents stimuli drawn from a large set of images<sup>1–8</sup> or sounds<sup>9,10</sup> and records one of two neuronal responses—the presence or absence of a spike. Psychophysical ‘classification image’<sup>11,12</sup> studies take a conceptually related approach. In this case, the response is a subject’s detection or lack of detection of a target embedded in experimentally controlled noise. The response being measured is different, but the goal of the analysis is similar—to determine which aspects of the stimuli lead to a particular neural or behavioral response (Fig. 1a). Functional imaging studies also share this basic design, but the high-dimensional variable is no longer under the experimenter’s control, so some aspects of the problem are reversed. For example, the investigator repeatedly presents stimuli from one of two categories and records many examples of images of neural activity elicited by the stimuli. In a simple optical imaging experiment<sup>13</sup> in

visual cortex, the two categories might consist of vertical and horizontal gratings; in a functional brain imaging experiment<sup>14</sup>, the two categories might consist of a behavioral task and a baseline state or two contrasting sets of sensory stimuli. In these experiments, the goal of the analysis is to determine which aspects of the brain image are associated with each category (Fig. 1b).

These are only the simplest prototypes. The multivariate quantity may be spatiotemporal sequences and not just static spatial images. The categorical quantity may have more than two possible values (for example, multiple orientations presented in an imaging experiment in visual cortex or temporal sequences of spikes in receptive field mapping). However, the essence of the analytic challenge in all of these experimental approaches remains one of relating a highly multivariate quantity to a categorical quantity. By far the most common, and perhaps most intuitive, strategy is fundamentally identical across these different experiments—it amounts to calculating the average value of the high-dimensional variable for each value of the categorical variable and subtracting one average from the other. This is also at the heart of reverse correlation techniques, which are pervasive in receptive field mapping or differential imaging.

However, intuitive methods may be inaccurate or inefficient and may overlook important details. In this article, I will consider this class of problems abstractly, highlighting certain similarities and differences between the problems faced by the different experimental approaches. My first goal is to illustrate the reasons why the intuitive analysis strategy may not be the best, and the conceptual challenges that must be faced. I will then describe (without technical detail) several approaches to these challenges that have been recently developed and applied within individual experimental approaches. I suggest that wider recognition of the common conceptual problem being solved in the different contexts, including both a focus on aspects that are specific to one approach as well as application of methods beyond their original context, will benefit both the development of analytic tools and the analysis of data.

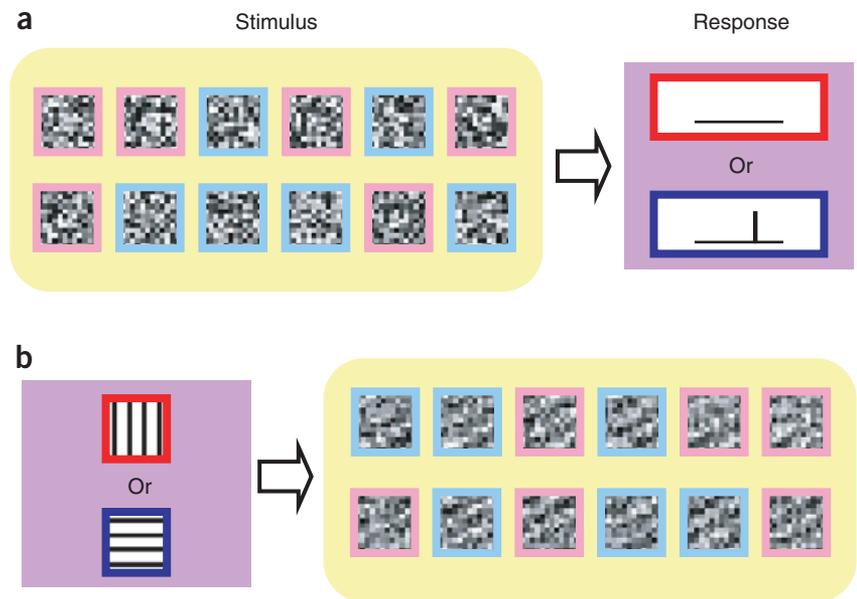
## Why analysis is challenging

To understand what analysis of these datasets might entail, consider a geometric view of a highly reduced experiment in which the goal is to categorize an image consisting of only two pixels (Fig. 2). In each panel, each image is represented by a point whose horizontal ( $x$ ) and vertical ( $y$ ) coordinates represent the image value of the two pixels. The color assigned to each point corresponds to the category associated with that image, as determined experimentally.

Jonathan D. Victor is in the Department of Neurology & Neuroscience, Weill Medical College of Cornell University, 1300 York Avenue, New York, New York 10021, USA.  
e-mail: jd victo@med.cornell.edu.

Published online 23 November 2005; doi:10.1038/nn1607

**Figure 1** Two kinds of experiments in which a highly multivariate quantity (yellow box) is to be related to a categorical quantity (purple box). (a) The multivariate quantity is the stimulus, and the categorical quantity is the response (as in receptive field mapping and classification image studies). (b) The multivariate quantity is the response, and the categorical quantity is the stimulus (as in optical imaging and functional brain imaging). In each case, the investigator determines the categories in advance (items outlined in red versus blue), but the instances of the multivariate quantity associated with each category (items outlined in pink versus pale blue) are determined from the experiment.



For definiteness, I will discuss the problem in terms of receptive field analysis, but the ideas apply equally well to classification image analysis. Consider an idealized (and highly simplified) receptive field mapping experiment (Fig. 2a) in which the stimulus consists of uncorrelated Gaussian noise at two pixels, and the neuron responds to only one pixel. The value of the stimulus in the pixel that drives the neuron is represented by the horizontal coordinate in Figure 2a. Because the pixel values in the stimuli are assumed to be uncorrelated, the stimuli form a circularly symmetric cloud. The neuron is assumed to respond to only one pixel, so the probability that a point is assigned to the two response categories depends only on its horizontal position. The demarcation between ‘blue’ responses (left half) and ‘red’ responses (right half) is not sharp, to represent the presence of neural noise that combines additively with the stimulus.

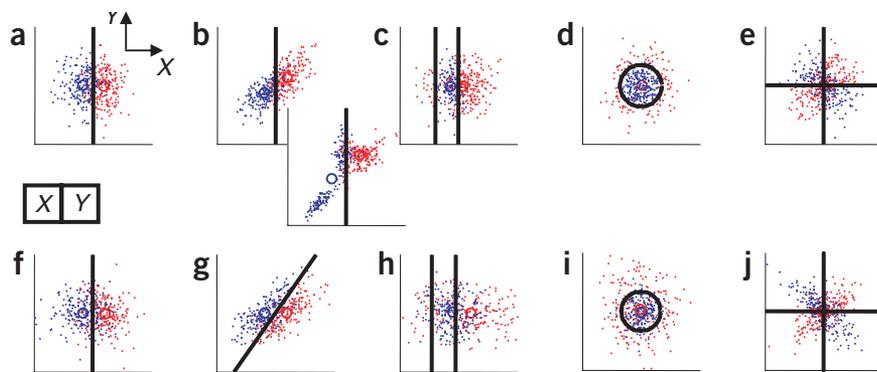
The red and blue circles represent the results of analyzing these data by determining the average stimulus that led to each response. This is the approach taken in the standard method of reverse correlation. The horizontal displacement of these circles properly reveal a dependence of the neural response on only one pixel value. Moreover, the best partitioning of the stimuli into subsets that are associated with the two responses (based on the experimental data) is given by the perpendicular bisector of the line between these averages. This line separates the points that are closest to the center (average) of the red cloud from those that are closest to the center of the blue cloud. Thus, averaging suffices to answer two basic questions: what is the average stimulus corresponding to each response ( $Q_{\text{center}}$ ), and what is the best way to determine which response will be elicited by a particular stimulus ( $Q_{\text{rule}}$ ). Averaging will answer these questions when the multivariate data are independent and identically distributed, the system is linear and noise is additive. As the next panels show, relaxation of these conditions can lead to very different results.

In cases where there is correlation between the two pixel values (Fig. 2b), averaging fails to capture the full relationship between stimulus and response. Such correlations are typically present in natural images. Even though the neuron only responds to the pixel represented by the horizontal coordinate, the average stimulus corresponding to each response class (circles) is now displaced along the diagonal, as a consequence of the pixel-to-pixel correlation present in the stimulus. Thus, the perpendicular bisector between the centers of the clouds, which is oblique, no longer represents the best rule for determining which response will be elicited by a given stimulus. Rather, the best rule remains a vertical line (as in Fig. 2a). In short, correlations within the stimulus set induce bias when averaging or reverse correlation is applied. This bias can be corrected if the stimulus correlations are known, and have

a sufficiently simple form (for example, if the correlation is Gaussian). However, for stimulus sets such as natural stimuli, these conditions do not apply<sup>6,7</sup>. When the correlation structure of the stimulus set is sufficiently complex, bias correction is problematic and the ‘average’ stimulus may not be typical of any stimuli (Fig. 2b, inset).

Even if the stimuli are uncorrelated, averaging will still be inadequate if the neuron is nonlinear. The next examples show how nonlinearity (non-additive combination of stimulus components) leads to additional complexities (Figs. 2c–e). Consider a neuron whose response depends on only one pixel, but this dependence has a small quadratic contribution in addition to the linear response (Fig. 2c). Because of this nonlinearity, stimuli with a large negative value at pixel  $x$  also lead to a response in the red category. More importantly, the optimal partitioning of the stimuli into classes corresponding to the two responses (that is, the description of what the two responses ‘represent’) consists of two vertical lines, not one. Averaging gives no hint of the bipartite distribution of red responses but rather misleadingly summarizes the distribution of red responses by a single point—which might even lie within the blue distribution. The latter situation would arise if the nonlinearity dominates, as would be the case in a stereotypical neuron with a symmetrical on-off response.

In another example of nonlinearity, the model neuron is an ‘energy’ unit<sup>15</sup> (Fig. 2d). It produces one response if  $x^2 + y^2$  (the energy) exceeds a criterion and the other response if it does not. The optimal partitioning of the stimuli into the two classes is a circle that lies on the energy threshold. Yet another example of nonlinearity is an idealized edge detector (Fig. 2e): this neuron produces one response if  $x$  and  $y$  have opposite sign and the other response if they do not—that is, its response is determined by an interaction of the two pixel values, the product  $xy$ . In this case, the optimal rule for partitioning stimuli is described by intersecting lines that run along the axes. For both types of neuron, however, the means of the stimulus subsets that correspond to the two response classes coincide (red and blue circles in Fig. 2d,e). This means that a correlation analysis will not detect any signal at all, even though there is simple (but nonlinear) relationship between stimuli and response. When there is a mixture of linear and nonlinear contributions (Fig. 2c), a correlation analysis properly indicates the neuron’s dependence on input stimuli, but the signal-to-noise (that is, the separation of the mean stimuli of each class) is less than that in Figure 2a because the average includes stimuli with large negative  $x$  values.



**Figure 2** A geometric view of associations of multivariate and categorical data. In each panel, each instance of the multivariate (here, bivariate) data is represented by a point whose coordinates  $X$  and  $Y$  are the values of its two components (in this case, pixel intensities). The color assigned to each point indicates which of the two categories is associated with it. (a–e) Experiments in which the multivariate quantity is the stimulus, and the categorical quantity is measured. (f–j) Experiments in which the multivariate quantity is measured, and the categorical quantity is the stimulus. a and f indicate the simplest situation (uncorrelated and Gaussian bivariate data, with category linearly determined by one of its components). Other panels introduce correlation structure into the bivariate data (b,g), nonlinearities (c,d,h,i), or both (e,j). The large colored symbols in each panel indicate the centroids of the respective clouds, and are superimposed in panels d, e, i and j.

For the converse problem, where the stimulus is the categorical variable and the multivariate quantity is measured (that is, in optical or functional imaging), analogous situations can arise (Figs. 2f–j). In the simple, ideal situation in which the stimulus activates only one pixel and there is uncorrelated, additive Gaussian measurement noise in both pixels (Fig. 2f), the cloud of responses elicited by each stimulus is circularly symmetric, and the horizontal displacement of these clouds represents the mean response. As described earlier (Fig. 2a), the mean response to each stimulus provides an unbiased estimate of the positions of these clouds, and the perpendicular bisector between these means is the optimal partitioning of the responses into two categories.

When a stimulus-driven response confined to one pixel adds to noise that is correlated across the two pixels (Fig. 2g)—for example, by vascular pulsations—the cloud of points that represents the pairs of pixel values corresponding to each stimulus becomes elongated and oblique because of the background correlation. The mean image elicited by each stimulus category is no different than that of Figure 2f. However, the optimal way to discriminate between these sets of images is no longer a vertical line: it is an oblique line whose slope is determined by the degree of correlation of the noise background. There is a second, more subtle, effect of the fact that the set of images elicited by each stimulus forms an elongated cloud. The oblique axis does not contribute to separation of the two clouds, but variability along it reduces the reliability of the estimates of the clouds' centers. Thus, noise correlation has two effects: the optimal rule for discriminating the two image classes does not correspond to the perpendicular bisector of the line between their means, and a more reliable estimate of the difference between the means can be obtained by eliminating dimensions that contain large variance and small signal.

As is the case for receptive field mapping, a nonlinear relationship between the stimulus and the multivariate response makes straightforward averaging inadequate for capturing the relationship between signal and the stimulus that drove it (Fig. 2h–j). There is evidence that brain states are manifest not only by mean activity but also by changes in power and correlation structure<sup>16–18</sup>, suggesting that such nonlinear relationships indeed exist. These nonlinear relationships affect the mean position of each cloud of points and the lines and curves that optimally separate

the clouds, in a way similar to the analogous nonlinear relationships shown in Figure 2c–e.

The challenges of analyzing a real dataset are substantially greater than these simple examples would suggest for several reasons. First, in a real dataset, variability would be much higher than in the illustration in Figure 2, so that the clouds would overlap to a much greater extent. Second, although we considered separately the effects of deviations from Gaussian uncorrelated noise (Fig. 2b,g), local nonlinearities (Fig. 2c,h), and spatial interactions (Fig. 2d,e,i,j), these phenomena are typically all present to some degree. Third, the dimensionality of the multivariate dataset is typically large: 100 to 1,000 in receptive field mapping or classification image experiments, and  $10^5$  or larger in optical or functional imaging experiments. Figure 2 considers only two-dimensional datasets. As a result, a typical dataset represents only a sparse sample of the multivariate distribution. Thus, in contrast to these examples where the dataset provides a good estimate

of the shape of the distribution of the multivariate quantity, good estimates of the distribution of the multivariate quantity may not be available.

These various considerations have led to the development of sophisticated methods for analysis of such datasets. However, development of methods to deal with the effects of correlations have generally proceeded along separate lines in receptive field mapping, image classification, and imaging contexts. Attempts to analyze nonlinearities have been developed primarily for the purpose of receptive field mapping, and not in the other contexts.

### Two basic questions

As highlighted in Figure 2, there are two basic questions that can be asked about the correspondence between a multivariate quantity and a categorical one:  $Q_{\text{center}}$ —what is the most typical value of the multivariate quantity that corresponds to each value of the categorical variable (that is, where is the center of each cloud)—and  $Q_{\text{rule}}$ —what is the best rule for distinguishing these clouds. I do not mean to imply that the answers to these questions are the endpoints of the analysis or suffice to draw scientific conclusions, merely that they are common conceptual places to begin.

For receptive field analysis, it is natural to focus on  $Q_{\text{rule}}$ . Even for a neuron with very simple properties, the center of the cloud will depend on the choice of stimuli used in an experiment. But the rule, which can be thought of as a computation performed on the stimuli and an indication of what a neuronal response represents, can in principle represent a more universal characterization of the neuron.

For imaging experiments,  $Q_{\text{center}}$  and  $Q_{\text{rule}}$  are interesting and quite distinct, even if the stimulus-response relationship is linear (Fig. 2g). This is because the multivariate quantities (the pixel values) are typically highly correlated, both by the underlying physiology and the physics of imaging. The center of each cloud indeed indicates the average response to each stimulus. However, the answer to  $Q_{\text{rule}}$  provides additional information—how best to 'read out' a pattern of activity. The distinction between the two questions cannot be avoided, because the experimenter cannot control the correlation structure of the multivariate data.

The answers to  $Q_{\text{center}}$  and  $Q_{\text{rule}}$  are usually displayed as maps, but it is important to emphasize that these maps represent very different things. For  $Q_{\text{center}}$ , the map is simply an instance of the multivariate quantity, but the answer to  $Q_{\text{rule}}$  is a rule, not an instance of the multivariate quantity. If the rule is linear, it can be rendered as a map in the following manner: a linear rule is characterized as an assignment of weights ('sensitivities') that multiply each pixel of the stimulus; the partitioning is based on a sum of these products. Thus, the map that portrays a linear rule is a map of sensitivities (that is, quantities whose units are the reciprocal of the units of the multivariate data). Nonlinear rules can also be displayed in a map-like fashion, but here, too, the map describes rules to be applied to stimuli, not stimuli themselves.

In the following sections, we consider a variety of approaches to answer  $Q_{\text{center}}$  and  $Q_{\text{rule}}$ .

### Uncorrelated multivariate data

Standard analysis (subtraction of mean responses for imaging<sup>13</sup>, cross-correlation for receptive field<sup>2,3,19</sup> or classification image<sup>11</sup> analysis), address  $Q_{\text{center}}$ .  $Q_{\text{center}}$  and  $Q_{\text{rule}}$  are equivalent only under very special circumstances (Fig. 2). When the linearity condition is relaxed but the multivariate quantity remains independent and identically distributed (Fig. 2c–e),  $Q_{\text{rule}}$  can be determined by generalizations of the cross-correlation approach. The desired characterization corresponds to the 'kernels' of Wiener-like<sup>2,3,19</sup> procedures. For example, the basic computation in the recently developed spike-triggered covariance method<sup>4,5,8</sup> is equivalent to the standard Lee-Schetzen cross-correlation estimate for the second-order kernel<sup>19,20</sup>, followed by diagonalization.

In settings such as receptive field or classification image analysis in which the investigator has control over the multivariate quantity, the use of 'designer' stimulus sets (sinusoidal sums<sup>21</sup> and m-sequences<sup>1</sup>) are particularly advantageous. Such approaches are often effective in characterizing nonlinear (as well as linear) response properties. This is because these finite stimulus sets are, in some sense, more nearly uncorrelated than a random sample drawn from a large uncorrelated ensemble.

However, 'designer' methods cannot be applied to imaging data or to situations in which natural scenes<sup>6,7,22</sup> or sounds<sup>10,23</sup> are used for receptive field determination. These multivariate stimulus sets typically contain strong correlations that cannot be controlled or fully characterized. Here, the answer to  $Q_{\text{rule}}$  provides information about the stimulus-response relationship that is not available from  $Q_{\text{center}}$ . Moreover, as will be sketched below,  $Q_{\text{rule}}$ , along with the covariance structure of the stimulus, can provide a better estimate of  $Q_{\text{center}}$  than averaging.

### Correlated multivariate data, linear relationship

Even when correlations within the high-dimensional variable are present, linear regression identifies the linear function of a stimulus sequence that does the best job (in the mean-squared sense) of predicting the binary response (spike versus no spike; target seen versus target not seen). Thus, it is a natural approach to finding  $Q_{\text{rule}}$  under the assumption that the stimulus-response relationship is linear and correlations are present within the high-dimensional variable. Also, it can be extended in a Wiener-like fashion to nonlinear relationships, as has been done in the context of receptive field analysis<sup>24</sup>. Of note, linear regression was used in the original description of the classification image method<sup>12</sup>. Other than a few exceptions<sup>25</sup>, this approach is not often taken because there are statistical difficulties that confound direct application of linear regression to the experimental contexts we are considering. However, as we next describe, there are recently developed techniques that can surmount these difficulties.

In essence, finding  $Q_{\text{rule}}$  via linear regression requires two steps (Supplementary Note online): (i) estimating the covariance matrix

$S$  of the multivariate set, and (ii) multiplication of the mean difference between the two multivariate clouds by  $S^{-1}$ , the matrix inverse of  $S$ . The covariance matrix  $S$  is a symmetric array whose entries  $s_{jk}$  are the correlations of the  $j$ th and  $k$ th pixels. For independent, identically distributed data,  $S$  is proportional to the identity matrix, and linear regression reduces to simple subtraction, as  $S^{-1}$  is also proportional to the identity. In imaging data, pixel values are coupled by motion, light scatter, blood flow and other physiologic factors; in receptive field analysis, pixel values are coupled by the statistics of natural scenes<sup>23,26</sup>. Thus, linear regression and the related approaches described below differ fundamentally from the subtraction method, as  $S$  is far from a multiple of the identity matrix.

The main pitfall in linear regression is that of overfitting: namely, the estimated answer to  $Q_{\text{rule}}$  may work well for the particular experimental sampling of the multivariate dataset but does not generalize to a larger sample. Linear regression provides an answer to  $Q_{\text{rule}}$  that generalizes if the correlation structure of the multivariate data is well characterized by the experimental sample. This characterization requires enough samples to estimate its covariance matrix  $S$ , along with an assumption, typically that the distribution of the multivariate dataset is Gaussian, to determine higher-order correlations from second-order correlations. For optical or functional imaging data, there are typically many more pixels ( $10^5$  to  $10^6$ ) than samples ( $\sim 10^4$ ). For receptive field or classification image data, the undersampling problem is present but less severe ( $10^3$  to  $10^4$  pixels; number of samples in the same range), but the correlation structure (especially of natural images) is likely to be very non-Gaussian. Because of undersampling and/or non-Gaussian characteristics of the multivariate data, direct application of standard linear regression is likely to produce results that are worse than simple subtraction.

However, extensions of linear regression<sup>27–29</sup> developed for imaging are applicable to the undersampled regime by recasting the problem in a form that does not require explicit inversion of  $S$ . These approaches focus on the eigenvectors of the covariance matrix  $S$  and related matrices (Supplementary Note). The eigenvectors of  $S$  (which are its 'principal components') are a small set of images from which all images can be reconstructed. Eigenvectors can be ranked in importance according to their eigenvalue. The larger the eigenvalue corresponding to an eigenvector, the greater the extent to which the stimulus set explores the corresponding image direction. On the basis of the eigenvalues, one can select a subset of eigenvectors within which a linear regression-like procedure can be carried out accurately. In the example of Figure 2g, such a procedure would correspond to restricting the estimation process to the direction that crosses the narrow axes of the ellipse and forgoing an attempt to estimate the position of the ellipse centers along their long axes, where variability is greater.

As detailed in the Supplementary Note, there are several useful variations on this theme. The 'truncated inverse' method<sup>28</sup> selects the eigenvectors of  $S$  whose eigenvalues are sufficiently large. More elaborate approaches select a subspace on the basis of not only on the overall covariance structure of  $S$  but also on the covariances within the subset of stimuli that lead to each categorical response. This includes the classic Fisher<sup>27</sup> discriminant method, which restricts analysis to the one-dimensional projection that optimally discriminates between Gaussian fits to the two clouds. The 'indicator function' method<sup>30</sup> projects into several dimensions (the 'canonical variates'), chosen on the basis of a significance criterion. Canonical variates are also the basis of a method for characterizing spatiotemporal aspects of images acquired in fast fMRI studies<sup>31,32</sup>. A further variation is the 'generalized indicator function method'<sup>29</sup>, which considers the eigenvectors of a linear mix of  $S$  and the within-group covariances and weights these eigenvectors in a graded fashion (Supplementary Note).

These approaches all represent methods of dimensionality reduction—selection of subspaces that are likely to contain a large signal-to-noise ratio (SNR). Within this general framework, independent component analysis<sup>33</sup> can be viewed as a strategy for using higher moments of the images to identify mixtures of subspaces that are likely to contain signal<sup>34</sup>. There are other approaches to enhance SNR in functional imaging data that exploit specific features of such images (such as vascular artifacts<sup>35</sup>), but because these approaches are domain specific, we do not discuss them here.

Dimensional reduction can also be viewed as a form of ‘regularization’<sup>36–38</sup> to avoid the pitfalls of estimating the covariance structure from very limited data. The generalized indicator function method<sup>29</sup> can be viewed as a regularized determination of canonical variates<sup>37,38</sup>. ‘Ridge regression’<sup>38</sup> (**Supplementary Note**) is a regularization strategy (widely applied outside of neuroscience) that chooses a compromise between the estimated covariance matrix and the identity matrix. Variations of ridge regression that take into account smoothness constraints have recently been used to extract receptive field maps from natural stimuli, both in the visual<sup>22</sup> and auditory<sup>10</sup> domains, but apparently have not been applied in optical or functional imaging.

### Correlated multivariate data, nonlinear relationship

In functional imaging, analytical methods (including the recent developments reviewed above) assume that the relationship between the imaged signal and neural activity is linear<sup>39</sup>, and the analytical focus has been on determining this relationship when the activity-dependent signal represents a very small fraction of the image. In receptive field analysis, it is generally considered that the neural response will be substantially greater than background when the stimulus is appropriate, but it is recognized that the stimulus-response relationship may not be linear<sup>6,40</sup>. This has driven the development of efficient methods for identifying nonlinear relationships that succeed even in the presence of strongly non-Gaussian multivariate data<sup>6,40</sup>. Moreover, the neural response, even if considered categorical, is generated in a manner that has stochastic and dynamic aspects. Understanding the implications of spike generation for the convergence and bias credentials of various estimation techniques<sup>40</sup>, and the interpretation of the resulting receptive fields<sup>41,42</sup> is another focus of current work.

### Time for a convergence?

One can readily identify several reasons for the separate development of analytical techniques in receptive field and classification image analysis on the one hand, and optical and functional imaging on the other, owing to various differences between these settings. However, I argue that the implications of these differences are less compelling than generally assumed and, consequently, that there are likely to be substantial opportunities for synergy.

Most obviously (**Fig. 1**), the categorical and multivariate characteristics of stimulus and response are swapped. This has two implications. If the experimenter is willing to choose a ‘designer’ stimulus set, then there are opportunities for improved experimental design in receptive field or classification image analysis that are not available for optical or functional imaging. We do not focus on these here, however, and these are irrelevant to receptive field and classification image studies using natural stimuli. The other implication is that this distinction leads to a difference in how noise is considered. In imaging, no threshold is typically postulated; rather, neuronal and measurement variability smoothly combine with the ‘signal’. For receptive field and classification image analysis, it is generally assumed that after a computation is performed on the multivariate quantity, there is a threshold (such as a firing threshold or a decision threshold) that may be in part stochastic.

However, many treatments seek a rule for the firing rate (or decision variable) that minimizes the mean-squared prediction error, rather than an explicit maximum-likelihood solution for a model with a threshold. Because a mean-squared error criterion is essentially a maximum-likelihood criterion for an assumed Gaussian noise, it is tantamount to ignoring the statistical consequences of the threshold.

Consequently, unless detailed dynamics of spike trains are of interest, the stimulus-response inversion does not prevent application of analysis methods for receptive field mapping to imaging, and vice-versa. Explicit modeling of spike train dynamics may result in further improvement and insight<sup>40–42</sup>. But dynamics (that is, the effect of stimulus sequence and time course on response) are also present in an optical or functional experiment, suggesting that techniques to examine dynamics developed for receptive field analysis might usefully be applied to imaging.

Another apparent difference between these settings is what is typically considered limiting. In classification image experiments, the number of trials that can be obtained may be limited to  $10^3$  to  $10^4$ , as each trial requires an explicit behavioral response. In imaging, the main hurdle to analysis is usually considered to be an intrinsically low SNR: 1 in  $10^3$  to 1 in  $10^4$  for optical imaging, and 1 in  $10^2$  for fMRI. There may also be sources of variability that lead to highly structured artifacts such as pulsatile movements of the tissue. Because receptive field studies often have the implicit goal of prediction of responses to stimuli outside of the experimental sample, undersampling of the stimulus space is often considered to be the main problem, as (even when the SNR is high) a far greater sampling is required to construct a nonlinear model than to construct a linear model. However, SNR, number of trials and sampling of the stimulus space are always limiting. Analytical approaches that make better use of a given data set to identify smaller signals, provide greater spatial detail or shorten the experiment time to obtain results of a given quality are always useful, especially given the increasing availability of computational resources and the costs (not just direct economic) of obtaining neurophysiological data. In sum, analyzing imaging data in the presence of structured noise is closely analogous to identification of receptive fields from ‘natural’ stimulus sets<sup>6</sup>, which have strong but incompletely determined statistical structure.

In optical and functional imaging, the relationship of the multivariate quantity to a behavioral index is typically assumed to be linear<sup>39</sup>, whereas a linear relationship is not always assumed in classification image analysis<sup>43,44</sup> and in receptive field mapping<sup>2,3</sup>. But there is increasing recognition<sup>16–18</sup> that stimulus-induced changes in the correlation structure of brain activity, and not just its mean level, are behaviorally and mechanistically relevant. Identifying such changes in imaging data is closely allied with identifying nonlinear aspects of a neural stimulus-response relationship in the receptive field mapping context.

Even for imaging studies that do not seek a nonlinear relationship between stimuli and the imaging data, inclusion of a nonlinearity in the model might nevertheless benefit the goal of signal extraction: that is, by identifying the rule that distinguishes responses to the members of the stimulus set. An analytic procedure that forces a nonlinear relationship to be modeled as a linear one necessarily causes some ‘signal’ (deviation of a systematic nonlinear response from a linear one) to appear as ‘noise’. Application of traditional nonlinear kernel methods is problematic, as the introduction of a large number of free parameters would likely defeat any benefit of capturing more signal. However, it would be very worthwhile to explore the use of recent receptive field mapping methods that seek simple nonlinear relationships in an efficient manner<sup>6,40</sup>. It should be emphasized that even if the relationship between the imaging signal and neural activity were strictly linear<sup>39</sup>, one would still expect this approach to be of value. Such linearity applies to the mean signal,

not its fluctuations, and in the noisy, multivariate regime, the latter may dominate the stimulus-response relationship. Moreover, mean neural activity may not be linearly related to the categorical variable—an issue that becomes relevant in experiments in which the categorical variable can take on more than two values, such as an orientation or spatial frequency experiment.

Conversely, to maximize the ability to test mechanistic or functional receptive field models, it is necessary to identify not only the large, readily resolvable components, but small contributions as well, as illustrated by the recent analysis of complex cell receptive fields in terms of spike-triggered covariances<sup>5,8</sup>. Moreover, it is evident that further insight into neuronal properties can be gleaned from receptive field characterization with stimulus sets have complex statistics, including natural scenes<sup>6,7,23</sup>. For both of these reasons, analytical methods developed for imaging that improve signal-to-noise by using subspace selection<sup>38</sup> may be useful. One might even envision that such methods could be further refined (for receptive field characterization) by guiding the estimation of covariances by the known statistical regularities of natural images<sup>26,45,46</sup>.

## Conclusion

As detailed above, the problem of identifying the relationship between a highly multivariate quantity and a categorical quantity (such as a discrete stimulus, a behavioral response or the presence of an action potential) is deceptively simple. It has been approached by a variety of analytical techniques, most often motivated by particular features of one of these contexts. We claim that the distinctions between experimental domains are not as deep as generally assumed, and we speculate that opportunities for progress will result from applying these techniques (or the ideas behind them) beyond their original domain.

The above considerations are only starting points, not an exhaustive list, and it is likely that the benefits will be relatively specific to particular experimental situations and goals. Whether cross-application of such methods and ideas will result in new qualitative insights, or merely incremental advances, is difficult to predict. Nevertheless, recognition of the close relationships between the mathematical challenges in these domains will enrich and accelerate the development of improved analytical techniques.

Note: Supplementary information is available on the Nature Neuroscience website.

## ACKNOWLEDGMENTS

The author thanks S. Klein, T. Yokoo, L. Paninski and P. Buzás for helpful discussions. This work is supported in part by grants EY7977 and EY9314 to J.D.V.

## COMPETING INTERESTS STATEMENT

The author declares that he has no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

- Sutter, E. in *Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks* (eds. Pinter, R. & Nabet, B.) 171–220 (CRC Press, Cleveland, 1992).
- Chichilnisky, E.J. A simple white noise analysis of neuronal light responses. *Network* **12**, 199–213 (2001).
- Marmarelis, P.Z. & Naka, K. White-noise analysis of a neuron chain: an application of the Wiener theory. *Science* **175**, 1276–1278 (1972).
- Simoncelli, E., Paninski, L., Pillow, J. & Schwartz, O. in *The Cognitive Neurosciences* 3rd edn. (ed. Gazzaniga, M.) (MIT Press, Cambridge, Massachusetts, 2004).
- Touryan, J., Lau, B. & Dan, Y. Isolation of relevant visual features from random stimuli for cortical complex cells. *J. Neurosci.* **22**, 10811–10818 (2002).
- Sharpee, T., Rust, N.C. & Bialek, W. Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput.* **16**, 223–250 (2004).
- David, S.V., Vinje, W.E. & Gallant, J.L. Natural stimulus statistics alter the receptive field structure of V1 neurons. *J. Neurosci.* **24**, 6991–7006 (2004).
- Rust, N.C., Schwartz, O., Movshon, J.A. & Simoncelli, E.P. Spatiotemporal elements of macaque V1 receptive fields. *Neuron* **46**, 945–956 (2005).
- Escabi, M.A. & Schreiner, C.E. Nonlinear spectrotemporal sound analysis by neurons in the auditory midbrain. *J. Neurosci.* **22**, 4114–4131 (2002).
- Machens, C.K., Wehr, M.S. & Zador, A.M. Linearity of cortical receptive fields measured with natural sounds. *J. Neurosci.* **24**, 1089–1100 (2004).
- Eckstein, M.P. & Ahumada, A.J., Jr. Classification images: a tool to analyze visual strategies. *J. Vis.* **2**, 1x (2002).
- Ahumada, A.J., Jr & Lovell, J. Stimulus features in signal detection. *J. Acoust. Soc. Am.* **49**, 1751–1756 (1971).
- Grinvald, A. Optical imaging of architecture and function in the living brain sheds new light on cortical mechanisms underlying visual perception. *Brain Topogr.* **5**, 71–75 (1992).
- Kwong, K.K. *et al.* Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci. USA* **89**, 5675–5679 (1992).
- Ohzawa, I., DeAngelis, G.C. & Freeman, R.D. Encoding of binocular disparity by complex cells in the cat's visual cortex. *J. Neurophysiol.* **77**, 2879–2909 (1997).
- Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P. & Andersen, R.A. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.* **5**, 805–811 (2002).
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A. & Arieli, A. Spontaneously emerging cortical representations of visual attributes. *Nature* **425**, 954–956 (2003).
- Rodriguez, E. *et al.* Perception's shadow: long-distance synchronization of human brain activity. *Nature* **397**, 430–433 (1999).
- Lee, Y. & Schetzen, M. Measurement of the kernels of a nonlinear system by cross-correlation. *Int. J. Control* **2**, 237–254 (1965).
- de Ruyter van Steveninck, R. & Bialek, W. Real-time performance of a movement sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc. R. Soc. Lond. B.* **234**, 379–414 (1988).
- Victor, J.D., Shapley, R.M. & Knight, B.W. Nonlinear analysis of cat retinal ganglion cells in the frequency domain. *Proc. Natl. Acad. Sci. USA* **74**, 3068–3072 (1977).
- Smyth, D., Willmore, B., Baker, G.E., Thompson, I.D. & Tolhurst, D.J. The receptive-field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *J. Neurosci.* **23**, 4746–4759 (2003).
- Theunissen, F.E. *et al.* Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* **12**, 289–316 (2001).
- Korenberg, M.J., Bruder, S.B. & McLroy, P.J. Exact orthogonal kernel estimation from finite data records: extending Wiener's identification of nonlinear systems. *Ann. Biomed. Eng.* **16**, 201–214 (1988).
- Levi, D.M. & Klein, S.A. Classification images for detection and position discrimination in the fovea and parafovea. *J. Vis.* **2**, 46–65 (2002).
- Simoncelli, E.P. & Olshausen, B.A. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
- Fisher, R.A. The use of multiple measurements in taxonomic problems. *Ann. Eugen.* **7**, 179–188 (1936).
- Gabbay, M., Brennan, C., Kaplan, E. & Sirovich, L. A principal components-based method for the detection of neuronal activity maps: application to optical imaging. *Neuroimage* **11**, 313–325 (2000).
- Yokoo, T., Knight, B.W. & Sirovich, L. An optimization approach to signal extraction from noisy multivariate data. *Neuroimage* **14**, 1309–1326 (2001).
- Everson, R., Knight, B.W. & Sirovich, L. Separating spatially distributed response to stimulation from background. I. Optical imaging. *Biol. Cybern.* **77**, 407–417 (1997).
- Friston, K.J., Frith, C.D., Frackowiak, R.S. & Turner, R. Characterizing dynamic brain responses with fMRI: a multivariate approach. *Neuroimage* **2**, 166–172 (1995).
- Worsley, K.J., Poline, J.B., Friston, K.J. & Evans, A.C. Characterizing the response of PET and fMRI data using multivariate linear models. *Neuroimage* **6**, 305–319 (1997).
- Bell, A.J. & Sejnowski, T.J. An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.* **7**, 1129–1159 (1995).
- Thomas, C.G., Harshman, R.A. & Menon, R.S. Noise reduction in BOLD-based fMRI using component analysis. *Neuroimage* **17**, 1521–1537 (2002).
- Carmona, R.A., Hwang, W.L. & Frostig, R.D. Wavelet analysis for brain function imaging. *IEEE Trans. Med. Imaging* **14**, 556–564 (1995).
- Tikhonov, A.N. & Arsenin, V.Y. *Solutions of Ill-Posed Problems* (Wiley, New York, 1977).
- Hastie, T., Buja, A. & Tibshirani, R. Penalized discriminant analysis. *Ann. Stat.* **23**, 73–102 (1995).
- Hastie, T., Tibshirani, R. & Friedman, J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer-Verlag, New York, 2001).
- Boynton, G.M., Engel, S.A., Glover, G.H. & Heeger, D.J. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* **16**, 4207–4221 (1996).
- Paninski, L. Convergence properties of three spike-triggered analysis techniques. *Network* **14**, 437–464 (2003).
- Aguera y Arcas, B. & Fairhall, A.L. What causes a neuron to spike? *Neural Comput.* **15**, 1789–1807 (2003).
- Pillow, J. & Simoncelli, E. Biases in white noise analysis due to non-Poisson spike generation. *Neurocomputing* **52–54**, 109–115 (2003).
- Neri, P. Estimation of nonlinear psychophysical kernels. *J. Vis.* **4**, 82–91 (2004).
- Neri, P. & Heeger, D.J. Spatiotemporal mechanisms for detecting and identifying image features in human vision. *Nat. Neurosci.* **5**, 812–816 (2002).
- Field, D.J. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* **4**, 2379–2394 (1987).
- Dong, D.W. & Atick, J.J. Statistics of natural time-varying images. *Netw. Comput. Neural Syst.* **6**, 345–358 (1995).

