

## Identifying Functional Bases for Multidimensional Neural Computations

**Joel Kaardal**

*jkaardal@physics.ucsd.edu*

**Jeffrey D. Fitzgerald**

*jfitzgerald@physics.ucsd.edu*

*Computational Neurobiology Laboratory and Crick-Jacobs Center for Theoretical and Computational Biology, Salk Center for Biological Studies, La Jolla, CA 92037; and Center for Theoretical Biological Physics, University of California, San Diego, La Jolla, CA 92037, U.S.A.*

**Michael J. Berry II**

*berry@princeton.edu*

*Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, U.S.A.*

**Tatyana O. Sharpee**

*sharpee@salk.edu*

*Computational Neurobiology Laboratory and Crick-Jacobs Center for Theoretical and Computational Biology, Salk Center for Biological Studies, La Jolla, CA 92037; and Center for Theoretical Biological Physics, University of California, San Diego, La Jolla, CA 92037, U.S.A.*

Current dimensionality-reduction methods can identify relevant subspaces for neural computations but do not favor one basis over the other within the relevant subspace. Finding the appropriate basis can simplify the description of the nonlinear computation with respect to the relevant variables, making it easier to elucidate the underlying neural computation and make hypotheses about the neural circuitry, giving rise to the observed responses. Part of the problem is that although some of the dimensionality reduction methods can identify many of the relevant dimensions, it is usually difficult to map out or interpret the nonlinear transformation with respect to more than a few relevant dimensions simultaneously without some simplifying assumptions. While recent approaches make it possible to create predictive models based on many relevant dimensions simultaneously, there still remains the need to relate such predictive models to the mechanistic descriptions of the operation of underlying neural circuitry. Here we demonstrate that transforming to a basis within the relevant subspace where the neural computation is best described by a given nonlinear function often makes it easier to interpret

**the computation and describe it with a small number of parameters. We refer to the corresponding basis as the functional basis, and illustrate the utility of such transformation in the context of logical OR and logical AND functions. We show that although dimensionality-reduction methods such as spike-triggered covariance are able to find a relevant subspace, they often produce dimensions that are difficult to interpret and do not correspond to a functional basis. The functional features can be found using a maximum likelihood approach. The results are illustrated using simulated neurons and recordings from retinal ganglion cells. The resulting features are uniquely defined and nonorthogonal, and they make it easier to relate computational and mechanistic models to each other.**

## 1 Introduction

---

A central challenge in the study of sensory processing neurons lies in discerning the features of a given stimulus that influence spiking behavior. It has become apparent in recent years that many neurons are selective for more than one stimulus feature (Atencio, Sharpee, & Schreiner, 2008; Cantrell, Cang, Troy, & Liu, 2010; Chen, Han, Poo, & Dan, 2007; Fairhall et al., 2006; Felsen, Touryan, Han, & Dan, 2005; Fox, Fairhall, & Daniel, 2010; Hong, Agüera y Arcas, & Fairhall, 2007; Horwitz, Chichilnisky, & Albright, 2005, 2007; Kim, Lazar, & Slutskiy, 2011; Maravall, Petersen, Fairhall, Arabzadeh, & Diamond, 2007; Rust, Schwartz, Movshon, & Simoncelli, 2005; Sincich, Horton, & Sharpee, 2009; Touryan, Lau, & Dan, 2002), sometimes as many as a dozen, leading to highly nonlinear and potentially very complex input-output relationships. Dimensionality-reduction methods, such as spike-triggered covariance (STC) and its extensions (Bialek & de Ruyter van Steveninck, 2005; de Ruyter van Steveninck & Bialek, 1988; Paninski, 2003; Park & Pillow, 2011; Pillow & Simoncelli, 2006; Schwartz, Chichilnisky, & Simoncelli, 2002; Schwartz, Pillow, Rust, & Simoncelli, 2006) or maximally informative dimensions (MID) and its extensions (Fitzgerald, Rowekamp, Sincich, & Sharpee, 2011; Rajan & Bialek, 2012; Sharpee, Rust, & Bialek, 2004), are able to produce linear subspaces that are relevant to the spiking activity of a neuron according to some metric (e.g., changes in stimulus covariance caused by spiking or mutual information). However, these and other dimensionality-reduction methods specify features not uniquely, but up to a linear combination of them. Often the results are presented in terms of orthogonal bases. Such orthogonal representations make it difficult to infer the corresponding neural computation that often involve a set of overlapping stimulus features, as in many types of motion computations (Gollisch & Meister, 2010). Here we show that for certain types of computations, the relevant subspace produced by these methods contains extractable information about the functional neural circuitry involved in sensory processing.

The STC or MID subspaces are defined by a set of basis vectors that are often interpreted as the stimulus features encoded by a neuron. However, it may be more appropriate to say that the features are defined by a functional basis, which spans the same subspace. The functional basis is a set of features that best accounts for the observed nonlinearities in the neural response using a predefined function. While the STC and MID bases certainly have a meaning in terms of their respective metrics, they might not coincide with the functional basis, which one may hope will yield insights into the underlying neural circuitry and may be amenable to interpretation. For example, the functional basis may be nonorthogonal, which cannot be matched by orthogonal STC bases. It is possible to order dimensions according to how much information they capture (Pillow & Simoncelli, 2006), but the resultant features also often come out to be orthogonal. In addition, the functional basis may help achieve a more concise description of the nonlinear gain function that describes how the neural spiking probability varies with stimulus components along the features. The nonlinear gain function, together with the set of stimulus features, forms the so-called linear-nonlinear (LN) model of the neural response (Chichilnisky, 2001; de Boer & Kuyper, 1968; Meister & Berry, 1999; Schwartz et al., 2006; Victor & Shapley, 1980).

Here we explore the Boolean operations corresponding to logical OR and logical AND computations, which are thought to describe the computations of some real neurons, such as those that contribute to translation invariance (Riesenhuber & Poggio, 1999; Serre, Wolf, Bileschi, Riesenhuber, & Poggio, 2007), various types of motion computation (Gollisch & Meister, 2010), build selectivity to increasingly complex stimulus features (Barlow & Levick, 1965), or implement coincidence detection (Carr & Konishi, 1988).

We show that operations of this kind have clearly defined functional bases. The functional bases may span the same subspace as STC basis, yet they allow easier interpretation of the neural computation and the likely neural architecture underlying it. We describe a way to extract the functional feature set from the STC basis and demonstrate the method on several biologically inspired stimulated neurons, as well as recordings of salamander retinal ganglion cells.

**1.1 Functional Bases.** Consider a neuron encoding a  $D$ -dimensional stimulus  $\mathbf{S}$ . The neural response  $y$  is binary (when considered in a small time window), with  $y = 0$  meaning the cell is silent and  $y = 1$  meaning the cell spikes. The instantaneous firing rate of such a neuron is proportional to  $P(y = 1|\mathbf{S})$ , the conditional probability of a spike. The neuron is said to be selective for  $n$  features of the stimulus if  $P(y = 1|\mathbf{S}) = f(x_1, x_2, \dots, x_n)$ , where  $x_i = \mathbf{S} \cdot \mathbf{v}_i$  is the projection of the stimulus onto a vector  $\mathbf{v}_i$ . The collection of vectors  $\{\mathbf{v}_i\}$  represents the stimulus features for which the neuron is selective.

*1.1.1 Nonorthogonal Feature Sets.* For neurons beyond the sensory transduction stage, the features  $\{\mathbf{v}_i\}$  are determined by the relevant stimulus features of afferent neurons and in general are not orthogonal. For example, a retinal ganglion cell may combine inputs from several bipolar cells that have overlapping receptive fields (Asari & Meister, 2012; Cohen & Sterling, 1991). The STC analysis would typically yield the sum and difference between the bipolar cells' receptive fields, as schematically illustrated in Figures 1 and 5. Although the neural response function  $f$  can be defined in terms of any linearly independent basis formed from a linear combination of the feature vectors, we are interested in finding a basis that is more closely related to its biological function. Many types of functions can be described using logical operations. Therefore, in this work, we focus on finding representations of  $f$  that make such logical descriptions explicit and refer to the corresponding basis as the functional basis.

As an example, consider a neuron selective for a set of features  $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ , using a noisy logical OR operation

$$f(x_1, x_2, \dots, x_n) = \begin{cases} 1, & \text{if } x_i + \eta_i \geq \theta_i \text{ for any } i \in \{1, 2, \dots, n\} \\ 0, & \text{otherwise} \end{cases}, \quad (1.1)$$

where  $\eta_i$  is gaussian noise added to  $x$  and  $\theta_i$  is the spiking threshold along dimension  $\mathbf{v}_i$ . A particular instance of this type of neuron is shown in Figure 1A for  $n = 2$ . In this figure, open circles represent stimuli for which  $y = 0$  and the filled circles represent stimuli for which  $y = 1$ , separated by the thresholds (dashed lines) along the dimensions  $\mathbf{v}_1$  and  $\mathbf{v}_2$ .

The horizontal  $s_1$  and vertical  $s_2$  axes of the plot form the basis that a method like STC or MID might recover (in this case, they actually correspond to STC dimensions), but a more natural choice of basis is given by  $\mathbf{v}_1$  and  $\mathbf{v}_2$ . These features  $\mathbf{v}_1$  and  $\mathbf{v}_2$  have a simple meaning in the computation performed by the cell (see equation 1.1). By comparison, the nonlinear gain function would require a more complex description in terms of  $s_1$  and  $s_2$ ; for example, the threshold value along the  $s_1$  dimension depends on the value of  $s_2$ . In a general case, the description of the nonlinear gain function along a set of  $n$  axes requires a large number of parameters, such as  $N_{\text{bins}}^n$  values, where  $N_{\text{bins}}$  is the number of bins along each of the dimensions.

This leads to the curse of dimensionality, which makes it difficult to map out or describe the nonlinear computation at a conceptual level with respect to more than a few relevant dimensions, although recent computational approaches (Park, Horwitz, & Pillow, 2011) make it possible to build predictive models with respect to many dimensions simultaneously. The transformation to the functional basis provides a way to model the nonlinear gain function with a small number of parameters, such as the threshold  $\theta_i$ , scale, and noise magnitude along each of the axes in the functional basis. The number of parameters increases linearly with the number of

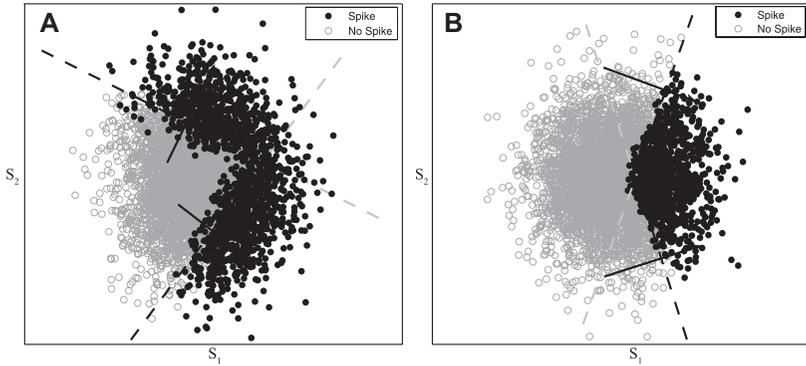


Figure 1: Computations in two dimensions with nonorthogonal features. Each dot represents a high-dimensional stimulus that has been projected onto this relevant subspace with axes  $s_1$  and  $s_2$  recovered by STC method. The filled circles are stimuli for which  $y = 1$  and the empty circles are those for which  $y = 0$ . Both logical OR (A) and logical AND operations (B) with two nonorthogonal features  $\mathbf{v}_1$  and  $\mathbf{v}_2$  are defined by two thresholds (dashed lines) that lie perpendicular to the feature directions. In the case of logical OR, any stimulus above either threshold causes a spike (corrupted by gaussian noise), whereas for a logical AND function, stimuli must be above all thresholds to elicit spikes. In both panels, the parts of the thresholds that determine spiking behavior are shown in black and the irrelevant parts in gray. Note that logical OR nonlinearity leads to the crescent-shape distribution of spike-eliciting inputs that are common in sensory systems (Fairhall et al., 2006).

relevant dimensions, and not exponentially as in the case of nonparametric descriptions.

Another computation relevant to biology is the logical AND operation, described by

$$f(x_1, x_2, \dots, x_n) = \begin{cases} 1, & \text{if } x_i + \eta_i \geq \theta_i \text{ for all } i \in \{1, 2, \dots, n\} \\ 0, & \text{otherwise} \end{cases} \quad (1.2)$$

An example of this function is shown in Figure 1B for  $n = 2$ . Again, the features defined by the threshold directions are not orthogonal and therefore could not be found using STC.

We note that the fact that the neural response function depends on  $n$  features (as in equations 1.1 and 1.2) does not exclude the possibility that the relevant subspace will have a lower dimensionality  $d < n$ . This will happen when the  $n$  features are linearly dependent. For instance, a Boolean function can have  $n$  planar thresholds in a  $d$ -dimensional relevant subspace where  $n > d$ . However in practice, it is unlikely that different inputs to

a neuron will have linearly dependent stimulus features. For example, if there are three presynaptic neurons, it is unlikely that each of them is not sensitive to at least partly novel aspects of the stimulus.

**1.2 Finding the Functional Bases for Boolean Computations.** The features encoded by the different pathways in a neural circuit are confined to a subspace in the high-dimensional stimulus space. Therefore, finding a functional basis can be simplified by first extracting this subspace using existing dimensionality-reduction methods such as STC. Once that subspace is determined, an analytical function representing the boolean OR and AND operations can be fit to identify a set of functional features.

*1.2.1 Dimensionality Reduction.* The STC method is designed to work with gaussian stimuli, and as such, we draw samples  $\mathbf{S}(t)$  for  $t = 1, 2, \dots, T$  from a  $D$ -dimensional zero mean normal distribution  $\mathcal{N}(\mathbf{0}, C_{\text{prior}})$ , where  $C_{\text{prior}}$  is the covariance matrix. Applying equation 1.1 with  $n$  features produces a model cell spike train  $\{y(t)\}$ , which can be used to calculate the spike-triggered covariance  $C_{\text{spike}}$ , defined by the elements

$$C_{\text{spike}}(i, j) = \frac{1}{T} \sum_{t=1}^T y(t) S_i(t) S_j(t) - \frac{1}{T^2} \sum_{t=1}^T y(t) S_i(t) \sum_{t=1}^T y(t) S_j(t). \quad (1.3)$$

The change in covariance caused by spiking is  $\Delta C = C_{\text{spike}} - C_{\text{prior}}$ . By diagonalizing  $\Delta C$ , we obtain the dimensions along which the covariance is altered by the neural response. These can be identified by the eigenvectors that have nonzero eigenvalues. This procedure is valid for uncorrelated or spherically symmetric stimuli (Bialek & de Ruyter van Steveninck, 2005; Chichilnisky, 2001; Samengo & Gollisch, 2013; Schwartz et al., 2006). Other methods can be used to find the relevant subspace for more complex stimuli (Fitzgerald, Rowekamp et al., 2011; Park & Pillow, 2011; Sharpee et al., 2004). Here we focus on uncorrelated stimuli, because this work is devoted to finding a specific basis within the relevant subspace, and not the subspace itself.

The accuracy with which the relevant subspace is found can be quantified using an overlap metric (Rowekamp & Sharpee, 2011) with respect to the subspace spanned by the features used by the model:

$$\mathcal{O} = \frac{\sqrt[d]{|\text{Det}(UV^T)|}}{\sqrt[2d]{|\text{Det}(UU^T)|} \sqrt[2d]{|\text{Det}(VV^T)|}}, \quad (1.4)$$

where  $U$  and  $V$  are  $d \times D$  matrices whose rows are the normalized eigenvectors  $\{\mathbf{u}_i\}$  and features  $\{\mathbf{v}_i\}$ , respectively. This metric is bounded between 0 and 1, with 0 meaning the two subspaces have no overlap and 1 meaning

they are identical, and gives us a way to evaluate our confidence in the dimensionality-reduction step. Note that if the number of features  $\{\mathbf{v}_i\}$  is larger than the dimensionality  $d$  of the relevant subspace, one can still choose  $d$  features that span that space to compute the overlap.

*1.2.2 Maximum Likelihood Models.* Given this reduced space, we wish to find the functional basis by looking for linear combinations of the vectors  $\mathbf{u}_i$  that best approximate a logical OR or logical AND computation. Such functions have the general form

$$P_{\text{OR}}(y = 1|\mathbf{S}) = 1 - \prod_{\alpha=1}^n [1 - g(U^T \mathbf{A}_\alpha \cdot \mathbf{S})], \quad (1.5)$$

$$P_{\text{AND}}(y = 1|\mathbf{S}) = \prod_{\alpha=1}^n g(U^T \mathbf{A}_\alpha \cdot \mathbf{S}), \quad (1.6)$$

where each  $\mathbf{A}_\alpha$  is a  $d$ -dimensional vector that describes the components of  $\alpha$ th axes of the functional basis within the STC basis. The function  $g(U^T \mathbf{A}_\alpha \cdot \mathbf{S})$  is some threshold-like function that describes the spike probability along the direction of the  $\alpha$ th functional feature of the computation.

To make progress, we must choose a specific functional form of  $g(U^T \mathbf{A}_\alpha \cdot \mathbf{S})$ . An obvious choice is the logistic function  $g_\alpha = 1/(1 + \exp(a_\alpha + U^T \mathbf{A}_\alpha \cdot \mathbf{S}))$ , where  $a_\alpha$  determines the location of the transition from silence to spiking. This choice of  $g$  is able to handle noise in the computation, as well as different thresholds along different features. It is also analytically quite tractable and has an information-theoretic interpretation as the least biased response function (Fitzgerald, Sincich, & Sharpee, 2011) consistent with encoding of a linear component along the functional feature.

The parameters  $\{\mathbf{A}_\alpha\}$  and  $\{a_\alpha\}$  can be found by maximizing the log likelihood. For a logical OR, this is

$$L_{\text{OR}} = \sum_{t=1}^N \left[ y(t) \log \left( 1 - \prod_{\alpha=1}^n \frac{1}{1 + e^{-a_\alpha - \mathbf{F}_\alpha \cdot \mathbf{S}(t)}} \right) - (1 - y(t)) \sum_{\alpha=1}^n \log(1 + e^{-a_\alpha - \mathbf{F}_\alpha \cdot \mathbf{S}(t)}) \right], \quad (1.7)$$

where  $\mathbf{F}_\alpha = U^T \mathbf{A}_\alpha$  are the functional basis vectors. Note that the functional basis vectors are not normalized to length 1. This optimization problem can be implemented using a straightforward conjugate gradient algorithm. In the end, the parameters  $\{a_\alpha\}$  that determine the threshold positions are unimportant if all that is desired is an estimate of the functional basis. The typical run time for the algorithm with 200,000 stimulus frames, each of

which had 256 pixels, is around 5 minutes on a 3.2 GHz desktop computer. Convergence time increases linearly with data set size and the number of basis vectors fit, and was not affected by the signal-to-noise ratio of the data set. The code is publicly available online at <http://cnl-t.salk.edu/Code/>.

## 2 Results

---

**2.1 Model Cells.** Three model cells were created to test whether maximizing equation 1.7 is able to recover the functional basis vectors. In all cases 200,000 stimulus samples were created from an uncorrelated gaussian distribution, with  $S_i \in \{0, \dots, 255\}$  representing the  $i$ th pixel intensity of an image. The response models were chosen to be logical OR functions, defined in the same manner as equation 1.1. The thresholds along each feature direction were set to the same value within a given model. The gaussian noise was added to the projections in such a way that the overall probability of a spike was between 0.2 and 0.4 for all models. The STC subspace in all cases had an overlap  $\mathcal{O} > 0.95$ , indicating a very good recovery of the model subspaces.

The first model considered is an example of a translationally invariant feature detector, shown in Figure 2A. The functional basis vectors are four shifted versions of a center-surround feature, located in the four corners of the image. These features overlap and are not orthogonal, and therefore the STC basis, shown in Figure 2B, is formed from linear combinations of these features. The logical OR functional dimensions make the computation performed by the cell difficult to interpret in the STC basis. The maximum likelihood OR solution for the functional basis with four features is shown in Figure 2C and matches quite well with the true features.

We also fit a maximum likelihood AND function to the STC basis for comparison. The outcome for the logical AND functional basis is shown in Figure 2D. Each of the four features is very nearly identical, indicating that the optimal solution for a logical AND is in fact not a logical AND at all, but rather a logistic function along a single dimension. This qualitative observation is quantified below.

The performance of these fits can be evaluated by comparing the predicted spike probability as a function of the stimulus to that obtained from the true model cell on repeated presentations of a group of stimuli. This is shown in Figure 2E for a 50-frame section of a 200,000-frame stimulus presented 10 times to the model cell. The spike probability obtained from the model (top panel) matches very well with the maximum likelihood logical OR (middle panel) and has a correlation coefficient of  $r = 0.97$  across all 200,000 frames. The maximum likelihood logical AND has a poorer performance, with  $r = 0.77$ . This comparison provides a way to determine how much like an OR or AND operation is the computation performed by a real cell.

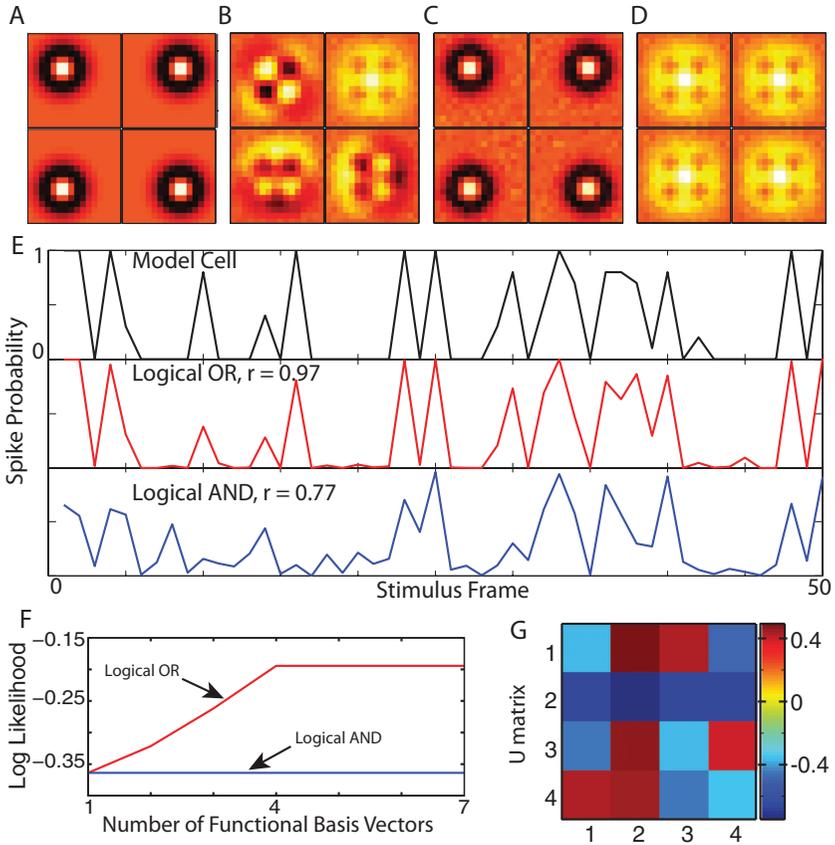


Figure 2: Translationally invariant model cell. The computation performed by the model is a logical OR operation on a center-surround feature, shifted to the four corners of the receptive field (A). The STC basis (B) and the maximum likelihood OR functional basis (C) both span the same space, but are much different in their appearance and interpretation. A maximum likelihood AND fit (D) finds four identical features, which do not match any of the model features. (E) Comparing the model spike probabilities generated from repeated presentations of a stimulus sequence to the predicted spike probabilities shows that the logical OR fit outperforms the logical AND fit significantly, with correlation coefficients of  $r = 0.97$  and  $r = 0.77$ , respectively. (F) The average log likelihood per time bin for logical AND and OR models with different number of functional basis vectors. (G) The transformation matrix from the STC to the functional bases (here, each vector was normalized to length 1 for presentation purposes).

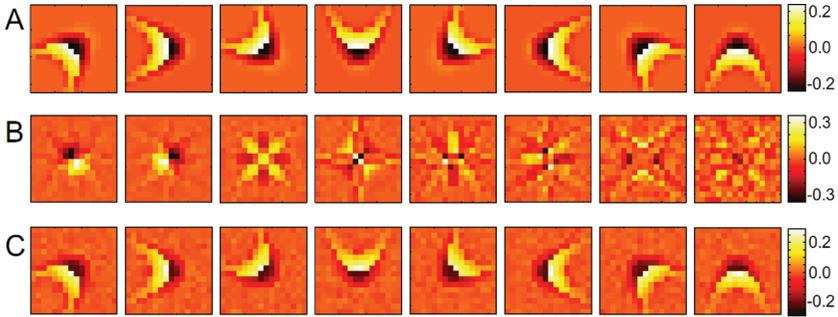


Figure 3: Rotationally invariant model cell. The computation performed by the model is a logical OR operation on a curved feature. The functional basis of the model (A) is defined by rotating the feature to eight different orientations. The (B) STC basis and (C) the maximum likelihood functional basis. The presentation order of the features in panel C was chosen to match the corresponding features of the model logical OR.

If one assumed no a priori knowledge of the number of features in this example, one can determine the optimal number of features to fit by performing a maximum log-likelihood calculation for a given number of features included in the model (see Figure 2F). In the case of the logical OR, the log likelihood increases as the features are added up to four. When the features are increased beyond four (the correct number for this model), no further increases in log likelihood are observed. This happens because extra features are either noise-like and have such large thresholds associated with them that they do not affect the spiking, or they reproduce one of the earlier features. In the case of the logical AND function, the log likelihood does not increase as more features are added to the model, because all of the features are roughly the same. Thus, finding the number of features where the log likelihood saturates provides a way to determine the correct number of features for a given model. The choice of a particular model structure, such as the logical AND or logical OR model, can be made according to which model yields a larger log likelihood.

As another example of an invariant computation, we tested a rotationally invariant model, shown in Figure 3. The model is selective for a curved feature at eight different orientations (see Figure 3A). Using STC, the eight-dimensional subspace was recovered (see Figure 3B) but had little resemblance to the functional feature set. The maximum likelihood estimate shown in Figure 3C, however, clearly recovers the functional basis.

Up to this point we have been fitting maximum likelihood functions with  $n_{\text{ML}} = n_{\text{model}}$  and considered the case with no overcompleteness where  $n_{\text{model}}$  equals the dimensionality  $d$  of the relevant subspace. However, the number of features is not known a priori, and therefore  $n_{\text{ML}}$  may be varied

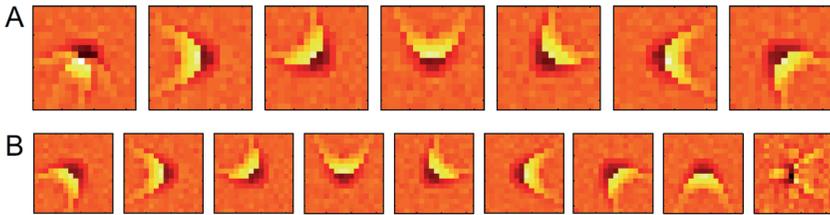


Figure 4: Changing the number of functional features. A maximum likelihood OR fit with  $n_{ML} = 7$  (A) and  $n_{ML} = 9$  (B). With the  $n_{ML} < n_{model}$ , some of the features appear as linear combinations of the true features, whereas for the  $n_{ML} > n_{model}$ , some of the features are redundant and begin to fit the noise in the computation. The threshold values along those features are so large that they seldom contribute to the model spiking.

to find the best fit. An example of this is shown in Figure 4 using the rotationally invariant model cell. The functional basis for  $n_{ML} = 7$ , Figure 4, recovers six of the features very well, while the seventh feature is a linear combination of the last two, with one orientation dominant over the other. For  $n_{ML} = 9$ , Figure 4B, recovers all eight features well, while the ninth is mostly redundant, with some slight overfitting to the noise in the neural computation. The redundant dimensions are associated with large threshold values such that they seldom contribute to the model spiking.

The value of the functional basis model in finding invariant dimensions is that no prior assumption is made about the type of invariance. By comparison, the recent method of maximally informative invariant dimensions (MIID) requires one to specify the type of invariance prior to doing the analysis (Eickenberg, Rowekamp, Kouh, & Sharpee, 2012; Vintch, Zaharia, Movshon, & Simoncelli, 2012). The MIID method also imposes a constraint that all of the features are identical except for shift in position or change in angle. Thus, functional basis and invariant methods approach the problem of characterizing neural feature selectivity with invariance from two different perspectives and in principle find the same dimensions if the assumptions made are congruent with the neural computation. Which method is more appropriate for analysis depends on what prior knowledge one has about the system and the degree to which invariance is exact.

To illustrate how the functional basis model can tackle a noninvariant case, we apply it to an example model cell with features that are not shifted or rotated copies of each other. Here, we considered computation inspired by properties of retinal ganglion cells (see Figure 5). This cell encodes nine blob-like features (see Figure 5A), which could represent the receptive fields of individual bipolar cells (Cohen & Sterling, 1991), and therefore the model could be interpreted as a logical OR retinal ganglion cell. Again the STC basis, shown in Figure 5C, makes it difficult to interpret the feature selectivity

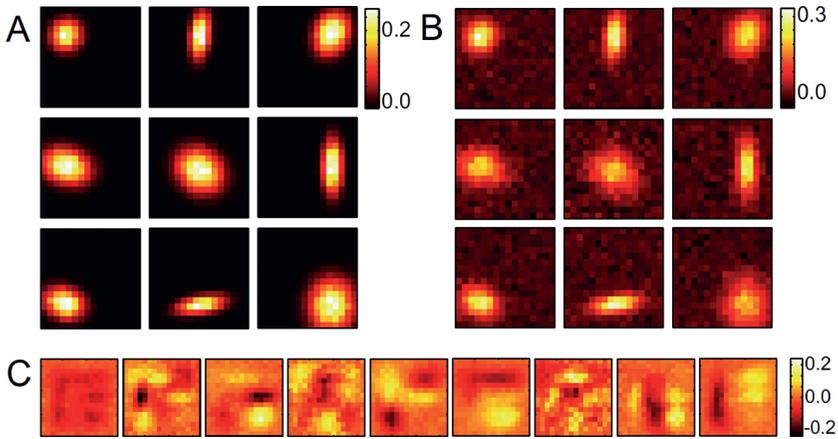


Figure 5: Logical OR model of a simulated retinal ganglion cell. Each of the nine blob-like features (A) represents the receptive fields of individual bipolar cells feeding into a retinal ganglion cell. (B) The maximum likelihood OR functional basis. (C) The STC basis.

of this model neuron. At the same time, the maximum likelihood OR linear combination of those vectors, shown in Figure 5B, matches well with the model features. Thus, one of the main advantages of using functional basis representation is that it can characterize the neural computation in cases where the invariance type is not known a priori or invariance is not exact.

**2.2 Application to Retinal Ganglion Cell Recordings.** We now illustrate how the approach works for characterizing encoding of retinal ganglion cells (RGCs) in the salamander retina probed with uncorrelated stimuli. The white noise stimulus consisted of binary checks forming frames with  $40 \times 40$  pixels. The stimulus had 137,145 frames at a rate of 60 Hz. There were 53 cells in the data set, with a wide range of activity from 118 to 30,908 total spikes (median 5453 spikes). The data set was collected as part of a previous study, and the electrophysiological methods are described there (Marre et al., 2012). As with the model cells, the first step of the analysis was to find the relevant subspace. We decided to focus on the spatial profile of the neurons as a way of illustrating the potential of this method for reconstructing receptive field properties of neurons (or parts of the neural circuit) that are presynaptic to ganglion cells. In the case of retinal ganglion cells, one might hope that the functional bases would then describe receptive fields of bipolar cells feeding into a given ganglion cell or nonlinear pieces within the dendritic computation of the ganglion cell (Soodak, Shapley, & Kaplan, 1991). To take into account temporal integration properties of RGCs, we computed the spatiotemporal spike-triggered average (STA)

for each neuron (de Boer & Kuypers, 1968). We then performed the singular value decomposition (Press, Teukolsky, Vetterling, & Flannery, 1992) of the spatiotemporal STA and took the first principal temporal vector as the estimate of the neuron's temporal kernel. Applying this filter to stimuli yields a set of spatial patterns that can then be associated with the measured neural responses. Subsequent analyses were carried out on  $16 \times 16$  patches of these temporally filtered spatial-only stimuli. The patches were centered at the location of the maximum of the STA of a given neuron. To find the relevant subspace, STC analysis was performed, and the relevant vectors were determined by comparing the resulting eigenvalue distribution with the eigenvalue distribution of a series of spike trains whose association with stimuli was broken. The association between stimuli and responses was broken by shifting the spike train forward by various amounts, with a minimal amount of the shift. The advantage of this procedure compared to shuffling spike trains is that it preserves all of the higher-order structure in the spike train in addition to the average spike rate (Bialek & de Ruyter van Steveninck, 2005). In our case, the minimal amount of shift was set to 100 frames. The STC analysis was then performed on 40 of such shifted spike trains, and the eigenvalue distributions from these analyses were combined. The eigenvalues from the measured (unshifted) spike train or stimuli pairs were considered significant if they lay outside the bounds of the eigenvalue distribution obtained with shifted spike trains. Of the 53 cells, 30 neurons had more than one relevant stimulus feature and therefore were of interest for further investigation using functional bases. The number of spikes available for these cells ranged between 192 and 21,894 (median 5426). For the example cell, which produced 8389 spikes, the STC method yielded four significant dimensions (see Figure 6A). The STC dimensions have commonly observed spatial profiles, with the dominant feature being uniphasic and the secondary, tertiary, and quaternary features containing two or more subregions of opposite polarity, similar to what was found in the primary visual cortex (Rust et al., 2005).

To find the functional basis for this example neuron, we fit both the logical OR and AND models. In this case, we focus the presentation of results on the logical OR model, which accounted better for the responses of this cell than the logical AND model did. First, the data set was split into four sections, as in the jackknife method (Efron & Tibshirani, 1998). By excluding one section for use as a test set to evaluate model performance while the other three sections were used as training sets to fit the model, we had four data sets with which to make a comparison of the functional models. Since the individual test sets have different information content, the log likelihoods cannot be directly compared. Rather, we computed a normalized log-likelihood difference to compare models A and B:

$$\Delta L_{A,B} = \frac{L_A - L_B}{|L_A| + |L_B|}, \quad (2.1)$$

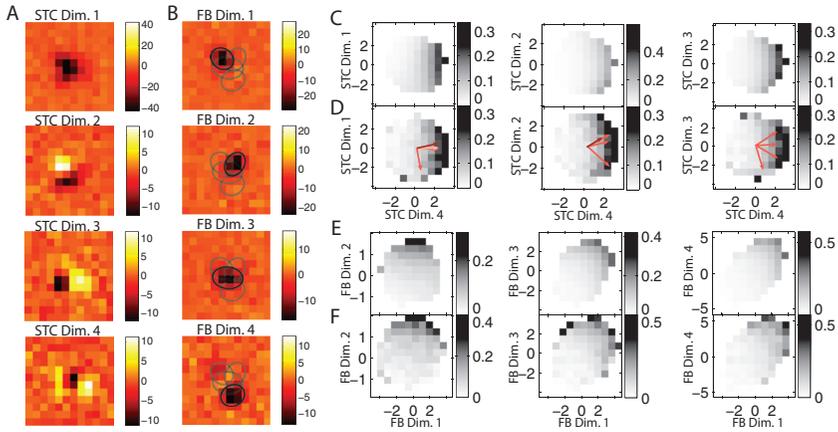


Figure 6: An example RGC stimulated by white noise. The STC basis (A) is rotated according to the logical OR functional basis method to produce the functional basis (B). The relevant image features are plotted relative to the average standard deviation per pixel. Unlike the STC basis, the image features in the functional basis are localized. To show the relative positioning of functional basis features relative to the neuron's receptive field, we fit these features with two-dimensional gaussians. Ellipses show one standard deviation contours from gaussian fits to the four functional basis features. Black ellipses correspond to the functional basis feature with which they are overlaid, and gray ellipses correspond to the other three functional basis features. We show the two-dimensional nonlinearity predicted (see equation 2.2) by the functional basis method (C), together with the empirically determined nonlinearity from the spiking data (D) in the STC basis. The predicted nonlinearity (E) also closely matches the empirical nonlinearity (F) when plotted in the functional basis projection space.

where  $L_A$  and  $L_B$  are the log likelihoods of model A and B, respectively, computed according to equation 1.7. If the mean log-likelihood difference across all test sets,  $\langle \Delta L_{A,B} \rangle$ , was significantly greater than 0, then model A was regarded as the better model. (The errors were computed using the jackknife method (Efron & Tibshirani, 1998) by scaling the standard error of the mean by  $t - 1$ , where  $t$  is the number of test data sets, to take into account that the training data sets are not independent.) For the example RGC neuron, the mean log-likelihood difference  $\Delta L_{\text{OR,AND}} = 0.0026 \pm 0.0001$ . Comparing logical OR and logical AND to the LN model with nonparametric nonlinearity yielded values  $\Delta L_{\text{OR,LN}} = 0.5454 \pm 0.006$  and  $\Delta L_{\text{AND,LN}} = 0.5436 \pm 0.006$ , respectively. This shows that for the example cell, logical OR and logical AND both perform better than the nonparametric LN model in making predictions, but the logical OR function is a better functional basis model than the logical AND. We note that increasing the number of

functional bases did not increase the log likelihood of the logical OR model (logical OR model with respect to five features yielded the log-likelihood difference  $\Delta L_{\text{OR}_5, \text{OR}_4} = -0.0042 \pm 0.0285$ ). The features of the functional basis corresponding to the logical OR model are shown in Figure 6B. Unlike the STC dimensions in Figure 6A, all of the components of the logical OR model in Figure 6B are spatially localized and describe sensitivity to decrements in light intensity in neighboring but displaced parts of the visual field. These dimensions may correspond to the receptive fields of bipolar cells feeding into this RGCs. The expected number of bipolar inputs is approximately 10 (Cohen & Sterling, 1991). Although this is somewhat greater than the four features we observe here, the number of bipolar inputs is known to vary depending on eccentricity, ganglion cell type, and species. It is also possible that the functional basis dimensions we observe here correspond to nonlinear subunits within the dendritic fields of RGCs, as has been discussed for cells in the cat (Soodak et al., 1991) where 2 to 4 subunits were observed. In any case, the subunits that we observe here are physiologically plausible and provide a way to describe functionally separate inputs to a given RGC.

One may also compare this approach to methods that characterize inputs by specifying that relevant features are the same up to position invariance, as in the MIID method. On one hand, the MIID method can allow one to potentially obtain more precise estimates of relevant features by pooling measurements from different parts of the visual field. On the other hand, one might be interested in the variability observed at different spatial positions. In particular, for the example cell we are considering, the four functional basis features are not exact copies of each shifted to different positions. To determine how important it was to consider the individuality of different functional basis dimensions, we compared the log likelihood of the functional basis logical OR model with that of models built on nine identical features shifted by one pixel from the peak of STA, respectively. The log-likelihood difference of the invariant dimension model and the log-likelihood logical OR model averaged across test sets for the example cell was  $\Delta L_{\text{OR}, \text{MIID}} = 0.8710 \pm 0.0011$ . (We also tried the translation-invariant model with larger shifts from the STA peak, but these yielded even worse predictive power.) Thus, in this case, the functional basis transformation yielded much better predictive power compared to the model where features are constrained to have the same profile shifted to different positions in the visual field.

Similar results were obtained across the population of 30 cells with multidimensional relevant subspaces. The number of relevant dimensions per neuron varied from two to four, with 15 cells having two features, 9 cells having three relevant features, and 6 cells having four relevant STC features. All but two of these cells were best fit by the logical OR functional basis model when compared to the logical AND model (see Figure 7A). No cell was better fit by the LN model than the functional basis model (see

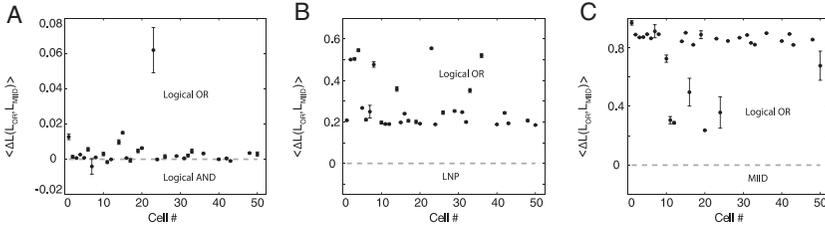


Figure 7: The difference in log likelihoods of the logical OR model against the logical OR model (A), LN model (B), and MIID (C) are plotted across the population of cells. The dashed line marks the threshold value at zero corresponding to the cross-over between the models that are being compared for the ability to predict the neural responses. The logical OR model outperformed the logical AND model and the nonparametric LN models for the majority of neurons.

Figure 7B). In fact, both functional basis models performed better than the nonparametric LN model for every cell. In all cases, the log likelihood saturated when the number of functional basis features chosen was equal to the number of STC basis vectors. Using the same procedure described in the previous paragraph, we applied MIID and compared it to the logical OR model. No cells were better described by invariant dimensions compared to functional basis representation (see Figure 7C).

The transformation to the functional basis can also be helpful for the interpretation of the nonlinear response function. For the example neuron in Figure 6D, the two-dimensional nonlinear firing rate function,

$$P(y = 1|x_i, x_j) = \sum_{k \neq i, j} \langle P(y = 1|x_1, x_2, \dots, x_d) \rangle_{x_k}, \quad (2.2)$$

with respect to STC dimensions  $s_i$  and  $s_j$  has a characteristic crescent-like shape, which is one of the types of nonlinearities previously reported in RGCs (Fairhall et al., 2006). The logical OR provides a natural way to explain the occurrence of crescent-like nonlinearities as arising from a sensitivity to similar but not identical features, each of which can trigger the neural response. Such nonlinearities are common in both visual (Fairhall et al., 2006; Rust et al., 2005) and auditory (Sharpee, Nagel, & Doupe, 2011) systems. It is noteworthy that the nonlinearity of the model neuron built according to the logical OR in Figure 1 also has a crescent shape. Returning to the example retinal neuron, we observe that the crescent shape of the nonlinearity observed in the STC basis (panels C and D) is unfolded by the transformation to the functional basis (panels E and F), as it should have been. Furthermore, the measured nonlinearities (panels D and F) are closely reproduced by the nonlinearities predicted by the model (panels C and E) in both the STC and functional bases.

### 3 Discussion

---

The concept of feature extraction using a linear subspace has found much success in the analysis of sensory processing neurons. However, simply discovering the relevant subspace is not necessarily the end of the story if the goal is to gain insight into the neural circuitry responsible for the processing of sensory information.

Dimensionality-reduction methods such as STC or MID identify the relevant subspace by finding a basis that captures significant changes in the stimulus covariance due to spiking or large amounts of the mutual information. Despite the meaning inherent in these bases, the resulting features are not necessarily the easiest to interpret in terms of functional computations. This is because standard dimensionality-reduction methods describe the neural input-output function using an orthogonal basis, whereas the relevant features of afferent inputs to a given neuron are often nonorthogonal. On the other hand, dimensionality-reduction methods for characterizing neural features that are explicitly based on assumptions of invariance (Eickenberg et al., 2012; Vintch et al., 2012) can account for nonorthogonal features but cannot capture individually in the receptive field subunits, also an important part of neural representation (Gauthier et al., 2009; Liu, Stevens, & Sharpee, 2009). The functional basis representation provides a compromise in that it does not impose a relationship between the relevant stimulus features, for example, they do not have to be shifted or scaled copies of each other. At the same time, searching for functional basis representations that correspond to biologically relevant nonlinearities, such as logical OR and AND nonlinearities (Riesenhuber & Poggio, 1999; Serre et al., 2007), allows one to determine types of invariance pertinent to the responses of a given neuron instead of imposing them a priori.

The results obtained by standard dimensionality-reduction techniques and the functional basis methods are related in that they share the same relevant stimulus subspace, differing only in the choice of basis within that subspace. However, finding a functional basis representation is not dependent on dimensionality-reduction methods such as STC or MID. In principle, a maximum likelihood model of a logical OR or AND response function with  $n$  features could be fit in the full  $D$ -dimensional stimulus space. The appropriate number  $n$  of significant features can be determined by maximizing the log likelihood with respect to the number of features just as was done in the reduced subspace. Performing the computation on the full stimulus space can be, however, a computationally daunting task. Thus, the search for the functional basis can be performed effectively and in a much more timely manner within the reduced subspace.

Although we have considered only two specific computational forms, the idea of finding a functional basis may be extended to other types of functions if simplifying features can be suitably defined. There are, however, many computations for which the functional basis is impossible to

define without detailed knowledge of the underlying circuitry. For instance, a radially symmetric function in two dimensions, such as in the energy model of a complex cell in area V1 of the visual cortex (Adelson & Bergen, 1985; Rust et al., 2005), is equally well described by any choice of basis. Even in this case, however, one can take advantage of the deviations from the true radial symmetry that characterize a particular complex cell under consideration to find the number of threshold-like units that can best account for its responses. Another example is when irregularities within the retinal circuitry can help separate individual inputs to a given RGC (Field et al., 2010; Liu et al., 2009; Soo, Schwartz, Sadeghi, & Berry, 2011). In summary, finding a functional basis can help bridge the divide between low-dimensional descriptions of neural responses and the underlying neural circuitry.

### Acknowledgments

---

We thank Ryan Rowekamp, Jonathan Aljadeff, and the rest of the Sharpee group for helpful conversations and the Aspen Center for Physics for their hospitality and facilitation of this collaboration. This research was supported by grants R01EY019493, K25MH068904, and R01EY014196 from the National Institutes of Health, grant 0712852 from the National Science Foundation, Alfred P. Sloan Research Fellowship, Searle Funds, the McKnight Scholarship, the Ray Thomas Edwards Career Development Award in Biomedical Sciences, and the W. M. Keck Foundation Research Excellence Award. Additional resources were provided by the Center for Theoretical Biological Physics (NSF PHY-0822283).

### References

---

- Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A. Opt.*, *2*, 283–299.
- Asari, H., & Meister, M. (2012). Divergence of visual channels in the inner retina. *Nat. Neurosci.*, *15*, 1581–1589.
- Atencio, C. A., Sharpee, T. O., & Schreiner, C. E. (2008). Cooperative nonlinearities in auditory cortical neurons. *Neuron*, *58*, 956–966.
- Barlow, H., & Levick, W. R. (1965). The mechanism of directionally selective units in rabbit's retina. *J. Physiol.*, *178*, 477–504.
- Bialek, W., & de Ruyter van Steveninck, R. R. (2005). *Features and dimensions: Motion estimation in fly vision*. q-bio/0505003.
- Cantrell, D. R., Cang, J., Troy, J. B., & Liu, X. (2010). Non-centered spike-triggered covariance analysis reveals neurotrophin-3 as a developmental regulator of receptive field properties of on-off retinal ganglion cells. *PLoS Comput. Biol.*, *6*, e1000967.
- Carr, C. E., & Konishi, M. (1988). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA*, *85*, 8311–8315.

- Chen, X., Han, F., Poo, M., & Dan, Y. (2007). Excitatory and suppressive receptive field subunits in awake monkey primary visual cortex (V1). *PNAS*, *104*, 19120–19125.
- Chichilnisky, E. J. (2001). A simple white noise analysis of neuronal light responses. *Network: Comput. Neural Syst*, *12*, 199–213.
- Cohen, E., & Sterling, P. (1991). Microcircuitry related to the receptive field center of the ON beta cell in the cat. *J. Neurophysiol.*, *65*, 352–359.
- de Boer, E., & Kuyper, P. (1968). Triggered correlation. *IEEE Trans. Biomed. Eng.*, *BME-15*, 169–179.
- de Ruyter van Steveninck, R. R., & Bialek, W. (1988). 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*, *265*, 259–265.
- Efron, B., & Tibshirani, R. J. (1998). *An introduction to the bootstrap*. London: Chapman and Hall.
- Eickenberg, M., Rowekamp, R. J., Kouh, M., & Sharpee, T. O. (2012). Characterizing responses of translation-invariant neurons to natural stimuli: Maximally informative invariant dimensions. *Neural Comput.*, *24*, 2384–2421.
- Fairhall, A. L., Burlingame, C. A., Narasimhan, R., Harris, R. A., Puchalla, J. L., & Berry II, M. J. (2006). Selectivity for multiple stimulus features in retinal ganglion cells. *J. Neurophysiol.*, *96*, 2724–2738.
- Felsen, G., Touryan, J., Han, F., & Dan, Y. (2005). Cortical sensitivity to visual features in natural scenes. *PLoS Biol.*, *e342*, 1819–1828.
- Field, G. D., Gauthier, J. L., Sher, A., Greschner, M., Machado, T., Shlens, J., et al. (2010). Functional connectivity in the retina at the resolution of photoreceptors. *Nature*, *467*, 673–677.
- Fitzgerald, J. D., Rowekamp, R. J., Sincich, L. C., & Sharpee, T. O. (2011). Second-order dimensionality reduction using minimum and maximum mutual information models. *PLoS Comput. Biol.*, *7*, e1002249.
- Fitzgerald, J. D., Sincich, L. C., & Sharpee, T. O. (2011). Minimal models of multidimensional computations. *PLoS Comput. Biol.*, *7*, e1001111.
- Fox, J. L., Fairhall, A. L., & Daniel, T. L. (2010). Encoding properties of Haltere neurons enable motion feature detection in a biological gyroscope. *Proc. Natl. Acad. Sci. USA*, *107*, 3840–3845.
- Gauthier, J. L., Field, G. D., Sher, A., Greschner, M., Shlens, J., Litke, A. M., et al. (2009). Receptive fields in primate retina are coordinated to sample visual space more uniformly. *PLoS Biology*, *7*(4), e1000063. doi:10.1371/journal.pbio.1000063
- Gollisch, T., & Meister, M. (2010). Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron*, *65*, 150–164.
- Hong, S., Agüera y Arcas, B., & Fairhall, A. L. (2007). Single neuron computation: From dynamical system to feature detector. *Neural Comput.*, *112*, 3133–3172.
- Horwitz, G. D., Chichilnisky, E. J., & Albright, T. D. (2005). Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. *J. Neurophysiol.*, *93*, 2263–2278.
- Horwitz, G. D., Chichilnisky, E. J., & Albright, T. D. (2007). Cone inputs to simple and complex cells in V1 of awake macaque. *J. Neurophysiol.*, *97*, 3070–3081.
- Kim, A. J., Lazar, A. A., & Slutskiy, Y. B. (2011). System identification of drosophila olfactory sensory neurons. *J. Computational Neuroscience*, *30*, 143–161.

- Liu, Y. S., Stevens, C. F., & Sharpee, T. O. (2009). Predictable irregularities in retinal receptive fields. *PNAS*, *106*, 16499–16504.
- Maravall, M., Petersen, R. S., Fairhall, A. L., Arabzadeh, E., & Diamond, M. E. (2007). Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol.*, *5*, 323–334.
- Marre, O., Amodei, D., Deshmukh, N., Sadeghi, K., Soo, F., Holy, T. E., et al. (2012). Mapping a complete neural population in the retina. *Journal of Neuroscience*, *32*, 14859–14873.
- Meister, M., & Berry II, M. J. (1999). The neural code of the retina. *Neuron*, *22*, 435–450.
- Paninski, L. (2003). Convergence properties of some spike-triggered analysis techniques. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems*, 15. Cambridge, MA: MIT Press.
- Park, I. M. M., & Pillow, J. W. (2011). Bayesian spike-triggered covariance analysis. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Weinberger (Eds.), *Advances in neural information processing systems*, 24 (pp. 1692–1700). MIT Press.
- Park, M., Horowitz, G. D., & Pillow, J. W. (2011). Active learning of neural response functions with gaussian processes. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Weinberger (Eds.), *Advances in neural information processing systems*, 24 (pp. 2043–2051).
- Pillow, J. W., & Simoncelli, E. P. (2006). Dimensionality reduction in neural models: An information-theoretic generalization of spike-triggered average and covariance analysis. *Journal of Vision*, *6*, 414–428.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in C: The art of scientific computing*. Cambridge: Cambridge University Press.
- Rajan, K., & Bialek, W. (2012). Maximally informative “stimulus energies” in the analysis of neural responses to natural signals. <http://arxiv.org/abs/1201.0321>
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.*, *2*, 1019–1025.
- Rowekamp, R. J., & Sharpee, T. O. (2011). Analyzing multicomponent receptive fields from neural responses to natural stimuli. *Network: Comput. Neural Syst.*, *22*, 45–73.
- Rust, N. C., Schwartz, O., Movshon, J. A., & Simoncelli, E. P. (2005). Spatiotemporal elements of macaque V1 receptive fields. *Neuron*, *46*, 945–956.
- Samengo, I., & Gollisch, T. (2013). Spike-triggered covariance: Geometric proof, symmetry properties, & extension beyond gaussian stimuli. *Journal of Computational Neuroscience*, *34*, 137–161.
- Schwartz, O., Chichilnisky, E. J., & Simoncelli, E. (2002). Characterizing neural gain control using spike-triggered covariance. In T. G. Dietterich, S. Becker, & Z. Ghahramani (Eds.), *Advances in neural information processing systems*, 14. Cambridge, MA: MIT Press.
- Schwartz, O., Pillow, J., Rust, N., & Simoncelli, E. P. (2006). Spike-triggered neural characterization. *Journal of Vision*, *6*, 484–507.
- Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., & Poggio, T. (2007). Robust object recognition with cortex-like mechanisms. *IEEE Trans Pattern Anal. Mach. Intell.*, *29*, 411–426.
- Sharpee, T. O., Nagel, K. I., & Doupe, A. J. (2011). Two-dimensional adaptation in the auditory forebrain. *J. Neurophysiol.*, *106*, 1841–1861.

- Sharpee, T., Rust, N., & Bialek, W. (2004). Analyzing neural responses to natural signals: Maximally informative dimensions. *Neural Computation*, *16*, 223–250.
- Sincich, L. C., Horton, J. C., & Sharpee, T. O. (2009). Preserving information in neural transmission. *J. Neurosci.*, *29*, 6207–6216.
- Soo, F. S., Schwartz, G. W., Sadeghi, K., & Berry II, M. J. (2011). Fine spatial information represented in a population of retinal ganglion cells. *J. Neurosci.*, *31*, 2145–2155.
- Soodak, R. E., Shapley, R. M., & Kaplan, E. (1991). Fine structure of receptive field centers of X and Y cells of the cat. *Visual Neuroscience*, *6*, 621–628.
- Touryan, J., Lau, B., & Dan, Y. (2002). Isolation of relevant visual features from random stimuli for cortical complex cells. *J. Neurosci.*, *22*, 10811–10818.
- Victor, J. D., & Shapley, R. (1980). A method of nonlinear analysis in the frequency domain. *Biophys. J.*, *29*, 456–483.
- Vintch, B., Zaharia, A., Movshon, J., & Simoncelli, E. P. (2012). Efficient and direct estimation of a neural subunit model for sensory coding. In P. Bartlett, F. Pereira, C. Burges, L. Bottou, & K. Weinberger (Eds.), *Advances in neural information processing systems*, *25* (pp. 3113–3121).

---

Received June 4, 2011; accepted January 29, 2013.