

Estimating Receptive Fields from Responses to Natural Stimuli with Asymmetric Intensity Distributions

Nicholas A. Lesica^{1,2✉*}, Toshiyuki Ishii^{1,3}, Garrett B. Stanley⁴, Toshihiko Hosoya¹

1 RIKEN Brain Science Institute, Wako-shi, Saitama, Japan, **2** School of Engineering & Applied Sciences, Harvard University, Cambridge, Massachusetts, United States of America, **3** Toho University, Funabashi-shi, Chiba, Japan, **4** Coulter Department of Biomedical Engineering, Georgia Institute of Technology & Emory University, Atlanta, Georgia, United States of America

Abstract

The reasons for using natural stimuli to study sensory function are quickly mounting, as recent studies have revealed important differences in neural responses to natural and artificial stimuli. However, natural stimuli typically contain strong correlations and are spherically asymmetric (i.e. stimulus intensities are not symmetrically distributed around the mean), and these statistical complexities can bias receptive field (RF) estimates when standard techniques such as spike-triggered averaging or reverse correlation are used. While a number of approaches have been developed to explicitly correct the bias due to stimulus correlations, there is no complementary technique to correct the bias due to stimulus asymmetries. Here, we develop a method for RF estimation that corrects reverse correlation RF estimates for the spherical asymmetries present in natural stimuli. Using simulated neural responses, we demonstrate how stimulus asymmetries can bias reverse-correlation RF estimates (even for uncorrelated stimuli) and illustrate how this bias can be removed by explicit correction. We demonstrate the utility of the asymmetry correction method under experimental conditions by estimating RFs from the responses of retinal ganglion cells to natural stimuli and using these RFs to predict responses to novel stimuli.

Citation: Lesica NA, Ishii T, Stanley GB, Hosoya T (2008) Estimating Receptive Fields from Responses to Natural Stimuli with Asymmetric Intensity Distributions. *PLoS ONE* 3(8): e3060. doi:10.1371/journal.pone.0003060

Editor: Bernhard Baune, James Cook University, Australia

Received: January 16, 2008; **Accepted:** July 9, 2008; **Published:** August 26, 2008

Copyright: © 2008 Lesica et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: NAL was supported by a Sigma Xi Grant-In-Aid of Research, a Post-doctoral Fellowship from the Japan Society for the Promotion of Science, and National Geospatial-Intelligence Agency Grant HM1582-05-C-0009. TI and TH were supported by the RIKEN Research Fund.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: lesica@zi.biologie.uni-muenchen.de

✉ Current address: Bernstein Center for Computational Neuroscience, Munich, Germany

Introduction

Traditionally, the response properties of sensory neurons have been studied using simple stimuli such as bars and sinusoidal gratings for vision, and clicks or pure tones for audition. More recently, the range of stimuli used to probe sensory function has been expanded to include more complex stimuli such as Gaussian white noise. While studies of responses to such artificial stimuli have provided the foundation for our understanding of sensory function, recent studies suggest that there may be fundamental differences between the neural responses to artificial stimuli and natural stimuli. For example, numerous studies have shown that natural stimuli are coded more efficiently than artificial stimuli in both the visual [1–4] and auditory [5–8] systems. Furthermore, there is evidence that models of sensory processing derived from responses to artificial stimuli are not sufficient to predict neural responses to natural stimuli [9–11]. These results suggest that if we hope to understand sensory function under natural conditions, we must study neural responses to natural stimuli directly.

Natural visual and auditory stimuli have complex statistical properties. For example, natural stimuli typically contain strong correlations, evidenced by power that decreases with increasing spatiotemporal or spectrotemporal modulation frequency as $1/f^\alpha$, with α typically between 1 and 3 [12–15]. Natural stimuli are also spherically asymmetric, meaning that the probability distribution of stimulus intensities is not symmetric about the mean intensity (in

contrast to, for example, Gaussian white noise) [13–17]. Unfortunately, these same complex statistical properties that differentiate natural stimuli from artificial stimuli also complicate the use of neural responses to natural stimuli in fitting models of sensory processing. With the most popular methods for characterizing sensory processing, reverse-correlation and spike-triggered averaging, an estimate of the linear filter or receptive field (RF) that provides the minimum mean squared error prediction of the neural response is computed as a weighted average of all stimuli, with each stimulus scaled by the magnitude of the response that it evoked. While these methods have proven extremely useful for characterizing the basic function of sensory systems (for a recent review, see [18]), they require that the stimulus is drawn from a spherically symmetric distribution in order to produce an unbiased RF estimate [18–22]. While this constraint may be satisfied by artificial stimuli such as Gaussian white noise, it is violated by the correlations and asymmetries typically found in natural stimuli, and, thus, under certain conditions, reverse correlation RF estimates computed from responses to natural stimuli can be biased.

A number of least-squares techniques in which the second-order stimulus correlations are essentially ‘divided out’ have been developed and used to estimate RFs from the responses of visual and auditory neurons to natural stimuli (for reviews, see [21,23–25]). In addition to correcting for the second-order correlations in the stimulus, these approaches also correct for asymmetries in the stimulus that are due to these correlations, but other asymmetries

that remain can bias the RF estimate. These effects were demonstrated in a recent simulation study that showed that even for a system consisting only of a cascade of a linear RF and a simple threshold nonlinearity, the interaction between higher-order correlations in the stimulus and the nonlinearity can lead to a biased RF estimate [26].

From an intuitive perspective, the bias in reverse correlation RF estimates caused by spherical asymmetries in the stimulus is similar to the error that would result from non-uniform stimulus sampling in a simple experiment. For example, in an attempt to characterize the ocular dominance of a neuron in the visual cortex based on the total number of spikes elicited by stimulation of each eye, it is clear that each eye must be stimulated the same number of times. If the number of stimuli presented to each eye is different, then the results must be explicitly corrected by dividing the total number of spikes elicited by stimulation of each eye by the number of times the eye was stimulated. A similar approach can be used to correct the bias in reverse correlation RF estimates spherical asymmetries in the stimulus.

Here, we develop a method for RF estimation from responses to natural stimuli that corrects for the biases introduced by spherical asymmetries by explicitly weighting the contribution of each stimulus to the RF estimate not only by the response it evokes, but also by its probability of occurrence relative to other stimuli with the same magnitude (vector norm). Through a series of simple examples using simulated neural responses, we illustrate how stimulus asymmetries can bias reverse correlation RF estimates (even for uncorrelated stimuli) and demonstrate how explicit correction for spherical asymmetries can remove this bias. We also demonstrate the application of the asymmetry correction method to experimental data by estimating the temporal RFs of retinal ganglion cells from responses to correlated, spherically asymmetric natural luminance sequences and using the RFs to predict responses to novel stimuli.

Analysis

In this section, we establish a linear-nonlinear (LN) response model and describe the conditions under which the reverse correlation technique provides an accurate RF estimate within this context. We show how stimulus correlations and asymmetries can bias reverse correlation RF estimates and detail a method for removing these biases.

A linear-nonlinear model for neural responses

We assume an LN model where the neural response is given by $r_i = f(s_i^T \mathbf{g})$, where r_i is the instantaneous firing rate of the neuron at time i , $f(\cdot)$ is a static nonlinear function,

$$\mathbf{s}_i = \begin{bmatrix} s_i \\ s_{i-1} \\ \vdots \\ s_{i-m+1} \end{bmatrix}$$

is the vector of the m most recent stimuli (we refer to m as the stimulus dimensionality), and

$$\mathbf{g} = \begin{bmatrix} g_1 \\ g_2 \\ \vdots \\ g_m \end{bmatrix}$$

is the time-invariant temporal RF. We assume that the stimulus is wide-sense stationary (i.e., its first- and second-order statistical properties are not changing over time) and has zero mean

$$\sum_{i=1}^{\infty} s_i = 0$$

A record of n stimulus/response observations can be summarized as:

$$\mathbf{r} = f(\mathbf{S}\mathbf{g}) \quad \text{where} \quad \mathbf{r} = \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_n \end{bmatrix} \quad \text{and} \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_1^T \\ \mathbf{s}_2^T \\ \vdots \\ \mathbf{s}_n^T \end{bmatrix} \quad (1)$$

are the vector of neural responses at each time and the matrix of stimulus vectors, respectively. Note that in the case where the input to the static nonlinear function $f(\cdot)$ is a vector, its output is also a vector of the same dimension.

Reverse correlation receptive field estimation

If the system were linear ($\mathbf{r} = \mathbf{S}\mathbf{g}$), the optimal estimate of the RF $\hat{\mathbf{g}}^*$ (i.e. that which minimizes the mean squared error between the actual response and that predicted by the RF estimate) would be given by

$$\hat{\mathbf{g}}^* = \arg \min_{\mathbf{g} \in \mathbb{R}^m} \|\mathbf{r} - \mathbf{S}\mathbf{g}\|$$

where $\|\cdot\|$ denotes vector norm (see equation 4). Many authors have shown that the solution to this equation is

$$\hat{\mathbf{g}}_{ls} = \mathbf{C}_s^{-1} \mathbf{S}^T \mathbf{r} \quad (2)$$

where $\mathbf{C}_s = \mathbf{S}^T \mathbf{S}$ is the autocovariance matrix of the stimulus (see, for example, [23]). For uncorrelated stimuli with \mathbf{C}_s proportional to the identity matrix, the reverse correlation estimate of the RF

$$\hat{\mathbf{g}} = \mathbf{S}^T \mathbf{r} \quad (3)$$

will be proportional to the optimal solution $\hat{\mathbf{g}}^*$. It has also been shown that this solution holds for the LN system described above, provided that the stimulus has zero mean and is drawn from a well sampled spherically symmetric distribution, i.e. all stimuli with the same ℓ^2 -norm

$$\|\mathbf{s}_i\| = \left(\sum_{j=0}^{m-1} s_{i-j}^2 \right)^{1/2} \quad (4)$$

must occur with the same probability

$$\forall \{\mathbf{s}_i, \mathbf{s}_j : \|\mathbf{s}_i\| = \|\mathbf{s}_j\|\} : P(\mathbf{s}_i) = P(\mathbf{s}_j) \quad (5)$$

and that the shape of the static NL is such that average stimulus evoking a response is non-zero (for proof, see [19]). Thus, for reverse correlation to provide an accurate RF estimate within the context of the LN model, the stimulus must be uncorrelated and spherically symmetric.

The utility of the reverse correlation approach is illustrated in the simple example shown in figure 1a. In this example, the stimulus is a Gaussian white noise sequence (representing, for example, luminance changes in a spatially uniform visual stimulus or the temporal modulations in a pure tone auditory stimulus). The Gaussian white noise stimulus is uncorrelated and spherically symmetric, as evidenced by the two-point stimulus intensity distribution and autocovariance matrix \mathbf{C}_S shown in figure 1a. We simulate responses to this stimulus using the LN model (equation 1), with $\mathbf{g} = [0.3, -0.15]$ (representing a weighted temporal summation of $m=2$ stimulus values) and $f(\cdot)$ a perfect half-wave rectifier ($f(x) = x$ for $x > 0$ and $f(x) = 0$ otherwise). Because the stimulus is uncorrelated and spherically symmetric, the reverse correlation estimate of the RF from the simulated responses $\hat{\mathbf{g}} = \mathbf{S}^T \mathbf{r}$ (red arrow) matches the actual RF (green arrow).

The effects of stimulus correlations

If the stimulus is correlated, then the reverse correlation estimate $\hat{\mathbf{g}}$ will be a biased version of the actual RF \mathbf{g} , with the bias determined by the autocovariance matrix \mathbf{C}_S . This is illustrated in the simple example shown in figure 1b. In this example, the stimulus is a Gaussian noise sequence with strong correlations, evidenced by the skewed two-point intensity distribution and the non-zero off-diagonal elements of the autocovariance matrix \mathbf{C}_S . Because of these correlations, the reverse correlation estimate of the RF $\hat{\mathbf{g}}$ (red arrow) computed from responses simulated with the LN model is biased toward the elongated dimension in the stimulus distribution and does not match the actual RF \mathbf{g} (green arrow).

Fortunately, the reverse correlation estimate $\hat{\mathbf{g}}$ can be modified to correct the bias due to the stimulus correlations. For example, the reverse correlation estimate can be multiplied by the inverse of the autocovariance matrix to produce the least-squares RF estimate $\hat{\mathbf{g}}_{ls}$ as described above. Indeed, this correction has been used to estimate RFs from responses to natural stimuli in many

brain areas [9,11,27–31]. An alternative approach is to transform the stimulus to remove correlations before computing the reverse correlation estimate [3,4]. For the LN model neuron described above we can transform the stimulus as follows:

$$\mathbf{r} = f(\mathbf{S}\mathbf{g}) = f(\mathbf{S}\mathbf{A}\mathbf{A}^{-1}\mathbf{g}) = f(\mathbf{\Sigma}\boldsymbol{\gamma}) \tag{6}$$

where $\mathbf{\Sigma} = \mathbf{S}\mathbf{A}$ and $\boldsymbol{\gamma} = \mathbf{A}^{-1}\mathbf{g}$. We can compute the reverse correlation estimate of the transformed stimulus $\mathbf{\Sigma}$ and the response \mathbf{r} , $\hat{\boldsymbol{\gamma}} = \mathbf{\Sigma}^T \mathbf{r}$, and obtain an estimate of the actual RF \mathbf{g} by inverting the transformation on the RF, $\hat{\mathbf{g}}_c = \mathbf{A}\hat{\boldsymbol{\gamma}}$. The matrix \mathbf{A} is chosen such that the transformed stimulus is uncorrelated (i.e. its autocovariance matrix is proportional to the identity matrix). The autocovariance matrix of the transformed stimulus \mathbf{C}_Σ can be written as:

$$\mathbf{C}_\Sigma = \mathbf{\Sigma}^T \mathbf{\Sigma} = (\mathbf{S}\mathbf{A})^T (\mathbf{S}\mathbf{A}) = \mathbf{A}^T \mathbf{S}^T \mathbf{S} \mathbf{A} = \mathbf{A}^T \mathbf{C}_S \mathbf{A}$$

Thus, we want to choose \mathbf{A} such that $\mathbf{A}^T \mathbf{C}_S \mathbf{A} = \mathbf{I}$. \mathbf{C}_S can be decomposed as $\mathbf{C}_S = \mathbf{V}\mathbf{D}\mathbf{V}^{-1}$ where \mathbf{V} is an orthogonal matrix of the eigenvectors of \mathbf{C}_S (i.e. the principal components of \mathbf{S}) with $\mathbf{V}^T = \mathbf{V}^{-1}$ and \mathbf{D} is a diagonal matrix with the corresponding eigenvalues $[\lambda_1^2, \lambda_2^2, \dots, \lambda_m^2]$ on the diagonal (or, alternatively, a similar representation can be obtained through singular value decomposition as described in [32]). If we choose $\mathbf{A} = \mathbf{V}(\mathbf{D}^{-1/2})^T$, then $\mathbf{C}_\Sigma = \mathbf{A}^T \mathbf{C}_S \mathbf{A} = \mathbf{D}^{-1/2} \mathbf{V}^T (\mathbf{V}\mathbf{D}\mathbf{V}^{-1}) \mathbf{V} (\mathbf{D}^{-1/2})^T = \mathbf{I}$ and the transformed reverse correlation estimate

$$\hat{\mathbf{g}}_c = \mathbf{A}\boldsymbol{\gamma} = \mathbf{A}\mathbf{\Sigma}^T \mathbf{r} \tag{7}$$

will not be biased by the stimulus correlations.

To demonstrate the utility of this correction, we return to the simple example of estimating a known RF from simulated responses to a correlated Gaussian noise stimulus shown in figure 1b. As described above, because of the correlations in the

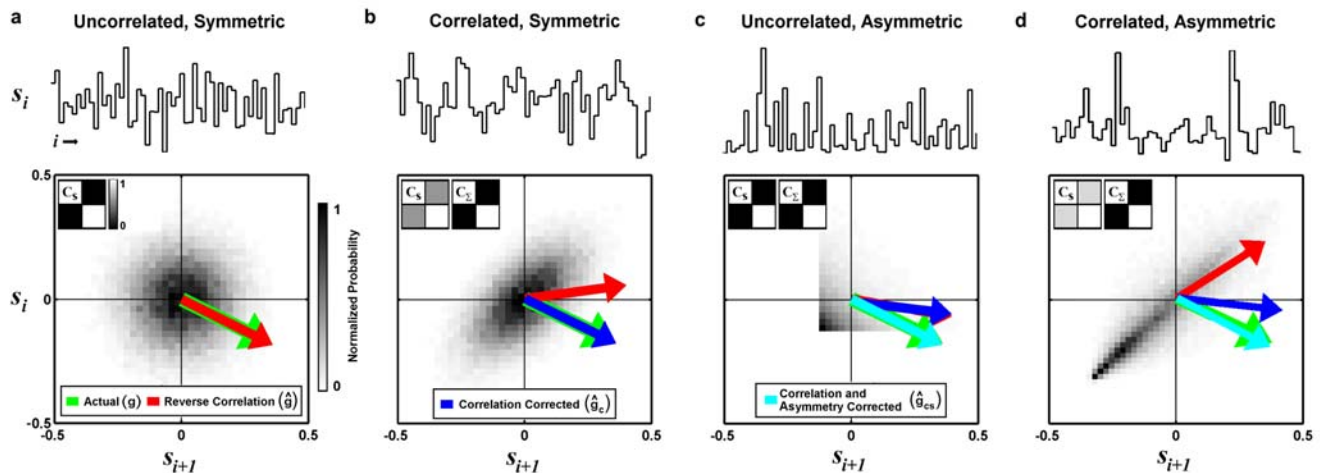


Figure 1. RF estimates from simulated responses to correlated and asymmetric stimuli. a) 60 sample segment of an uncorrelated and symmetric Gaussian white noise stimulus (top) along with its two-point intensity distribution (bottom) and autocovariance matrix \mathbf{C}_S (bottom inset). The actual RF $\mathbf{g} = [0.3, -0.15]$ ($m=2$) used to simulate LN model responses to the stimulus (green arrow) is shown along with the reverse correlation estimate of the RF $\hat{\mathbf{g}}$ computed from the simulated responses (red arrow). b) Results for a correlated Gaussian noise stimulus, presented as in a, along with the autocovariance matrix of the transformed stimulus \mathbf{C}_Σ and the reverse correlation RF estimate after correction for stimulus correlations $\hat{\mathbf{g}}_c$ (blue arrow). c) Results for an uncorrelated stimulus drawn from an exponential distribution, presented as in b, along with the reverse correlation RF estimate after correction for stimulus correlations and asymmetries $\hat{\mathbf{g}}_{cs}$ (cyan arrow). d) Results for a correlated and asymmetric stimulus, a time-series of natural intensities taken from the database of van Hateren [16], presented as in c. doi:10.1371/journal.pone.0003060.g001

stimulus, the reverse correlation estimate $\hat{\mathbf{g}}$ is biased. However, if we apply the transformation described in equation 6 to correct for the stimulus correlations, then the autocovariance matrix of the transformed stimulus \mathbf{C}_Σ has zero off-diagonal elements and the resulting RF estimate $\hat{\mathbf{g}}_c$ (blue arrow) now matches the actual RF \mathbf{g} .

Note that the transformed reverse correlation estimate $\hat{\mathbf{g}}_c = \mathbf{A} \Sigma^T \mathbf{r}$ can also be written as $\hat{\mathbf{g}}_c = \mathbf{A} \mathbf{A}^T \mathbf{S}^T \mathbf{r} = \mathbf{V} (\mathbf{D}^{-1/2})^T \mathbf{D}^{-1/2} \mathbf{V}^T \mathbf{S}^T \mathbf{r}$. In this form, it is clear that in computing the transformed reverse correlation estimate, each principal component of \mathbf{S} is multiplied by a factor related to the inverse of the corresponding eigenvalue. Thus, those principal components with smallest eigenvalues will have the largest effect on the estimate. For high dimensional natural stimuli with strong correlations, the difference between the largest and smallest eigenvalues can be several orders of magnitude (i.e. the condition number of the stimulus autocovariance matrix can be extremely large), and the effect of the principal component with largest eigenvalue on the RF estimate can be dwarfed by that of the principal component with the smallest eigenvalue. In this case, the RF estimate will be largely determined by those principal components along which the stimulus has the smallest variance and, under experimental conditions where only a limited number of noisy responses are observed, a large difference in eigenvalues can result in an RF estimate that is dominated by noise.

A number of approaches have been proposed to address this problem. For example, Theunissen and colleagues [9,23] have computed the transformed RF estimate using only those principal components with eigenvalues larger than some threshold value (i.e. if an eigenvalue is less than the threshold value, then the element of $\mathbf{D}^{-1/2}$ corresponding to that eigenvalue is set to zero). We adopted a variant of this approach in which only those principal components required to explain a certain fraction $0 < \epsilon \leq 1$ of the variance in the stimulus were retained. In addition to reducing the noise in the transformed reverse correlation RF estimate, eliminating some fraction of the principal components allows the stimulus to be represented in a lower dimensional space $m^* \leq m$, which simplifies the estimation of the stimulus probability distribution $P(\mathbf{s})$ as described below.

The effects of stimulus asymmetries

The transformation described above corrects the bias in reverse correlation RF estimates due to the second-order correlations in the stimulus, as well as the bias due to any asymmetries in the stimulus intensity distribution that result from those correlations. However, complex stimuli can contain additional asymmetries and these asymmetries can also bias RF estimates. This is illustrated in the simple example shown in figure 1c. In this example, each value of the stimulus is drawn from an uncorrelated exponential distribution, and there are clear asymmetries in the two-point intensity distribution. Because of these asymmetries, the reverse correlation estimate of the RF $\hat{\mathbf{g}}$ (red arrow) computed from simulated response of the LN model is biased and does not match the actual RF (green arrow). Additionally, because the stimulus autocovariance matrix $\mathbf{C}_\mathbf{s}$ is already proportional to the identity matrix, applying the transformation described in equation 6 to correct for the stimulus correlations has no effect, and the transformed reverse correlation RF estimate $\hat{\mathbf{g}}_c$ (blue arrow) is also biased.

Fortunately, the bias in the RF estimate due to stimulus asymmetries can also be corrected. Conceptually, the correction necessary to remove the bias due to stimulus asymmetries is analogous to the transformation used to correct the bias due to stimulus correlations. In removing the bias due to stimulus correlations, each principal component is weighted by the inverse

of the amount of stimulus variance that it explains, such that the effective contribution of every principal component to the transformed stimulus is the same. Similarly, the bias due to stimulus asymmetries can be removed by weighting each stimulus by its probability of occurrence relative to those of other stimuli with the same vector norm, such that the effective probability distribution of the stimulus is spherically symmetric. If we want to correct the biases due to stimulus correlations and asymmetries simultaneously, then we can estimate the RF as:

$$\hat{\mathbf{g}}_{cs} = \mathbf{A} \tilde{\Sigma}^T \mathbf{r}, \quad \text{where} \quad \tilde{\Sigma} = \begin{bmatrix} \tilde{P}(\sigma_1) \sigma_1^T \\ \tilde{P}(\sigma_2) \sigma_2^T \\ \vdots \\ \tilde{P}(\sigma_n) \sigma_n^T \end{bmatrix} \quad (8)$$

where $\tilde{P}(\sigma_i) = \tilde{P}(\|\sigma_i\|) P(\sigma_i)^{-1}$ is the asymmetry correction for a particular (transformed) stimulus and $\tilde{P}(\|\sigma_i\|)$ is the mean probability of occurrence of all stimuli with same vector norm as stimulus σ_i .

The asymmetry correction described in equation 8 requires estimation of the overall probability distribution of the stimulus $P(\sigma)$ and norm-specific probabilities $\tilde{P}(\|\sigma_i\|)$. In practice, we estimate $P(\sigma)$ by grouping stimuli into evenly spaced bins that span the range of stimulus values. $P(\sigma)$ is estimated not from the original stimulus, but from the coefficients that define each stimulus within the space defined by the principal components in \mathbf{V} , thus reducing the dimensionality of $P(\sigma)$ from m to m^* (note that for natural stimuli, which typically contain strong correlations, the value of m^* that results in the best RF estimates is often much less than m). The norm-specific probabilities $\tilde{P}(\|\sigma_i\|)$ are estimated by dividing the range of norms into evenly spaced bins and taking the mean of the probabilities of all stimuli whose norms fall within each bin. In general, we found that the effect of the correction for stimulus asymmetries was robust to changes in the number of bins used to estimate $P(\sigma)$ and $\tilde{P}(\|\sigma_i\|)$, even for relatively high-dimensional problems. Thus, for all of the simulated and experimental examples below, the probability distribution of the stimulus $P(\sigma)$ was computed after dividing the range of stimulus values into 250 evenly spaced bins and the norm-specific probabilities $\tilde{P}(\|\sigma_i\|)$ were computed after dividing the range of norms into 250 evenly spaced bins.

To demonstrate the utility of the asymmetry correction described in equation 8, we return to the simple example of estimating a known RF from simulated responses to a stimulus drawn from an uncorrelated exponential distribution shown in figure 1c. As described above, because of the asymmetries in the stimulus, both the reverse correlation estimate $\hat{\mathbf{g}}$ and the transformed reverse correlation estimate $\hat{\mathbf{g}}_c$ are biased. However, if we apply the asymmetry correction described in equation 8, then the resulting RF estimate $\hat{\mathbf{g}}_{cs} = \mathbf{A} \tilde{\Sigma}^T \mathbf{r}$ (cyan arrow) now matches the actual RF (green arrow).

Temporal receptive field estimation from simulated responses to natural stimuli

A final simple example shown in figure 1d demonstrates the utility of the asymmetry correction for a stimulus that is both correlated and asymmetric. In this example, the stimulus is a time-series of natural intensities taken from the database of van Hateren [16]. This stimulus contains strong correlations and asymmetries, as illustrated in two-point intensity distribution and stimulus autocovariance matrix $\mathbf{C}_\mathbf{s}$ shown in figure 1d. Because of these correlations and asymmetries, the reverse correlation estimate of the RF $\hat{\mathbf{g}}$ (red arrow) computed from simulated responses of the

LN model is biased and does not match the actual RF \mathbf{g} (green arrow). Transformation of the stimulus to remove the bias due to correlations improves the estimate $\hat{\mathbf{g}}_c$ (blue arrow), but only after correction for both stimulus correlations and asymmetries does the estimate $\hat{\mathbf{g}}_{cs}$ (cyan arrow) match the actual RF.

The simple examples in figure 1 demonstrate the ability of the asymmetry correction to improve estimates of low-dimensional ($m=2$) RFs. However, the RFs estimated from experimental responses of sensory neurons data typically have a much higher dimensionality ($\sim 10 < m < \sim 1000$). Because the asymmetry correction depends on the estimation of the m^* -dimensional probability distribution of the stimulus $P(\boldsymbol{\sigma})$, the efficacy of the correction may decrease as m increases.

To improve the efficacy of the asymmetry correction for higher-dimensional RFs, we added two additional parameters. The first parameter, θ , specifies the vector norm threshold that determines whether stimuli and their corresponding responses are included or excluded from the RF estimate (only stimuli with vector norms below θ are included in the estimate). The second parameter, ϕ , specifies the maximum value for the asymmetry correction $\tilde{P}(\boldsymbol{\sigma}_i)$. After rescaling the asymmetry corrections for all stimuli such that $1 \leq \tilde{P}(\boldsymbol{\sigma}_i) < \infty$, any $\tilde{P}(\boldsymbol{\sigma}_i) > \phi$ are set equal to ϕ . For natural stimuli, the probability of a stimulus $P(\boldsymbol{\sigma})$ tends to decrease with increasing norm and excluding those low probability stimuli with large norms or limiting the value of their asymmetry correction $\tilde{P}(\boldsymbol{\sigma}_i)$ can improve the efficacy of the overall asymmetry correction.

Including the two new parameters θ and ϕ , the asymmetry correction can be written as

$$\hat{\mathbf{g}}_{cs} = \mathbf{A}\boldsymbol{\Sigma}^{*T} \mathbf{r}, \quad \text{where} \quad \boldsymbol{\Sigma}^* = \begin{bmatrix} P^*(\boldsymbol{\sigma}_1)\boldsymbol{\sigma}_1^T \\ P^*(\boldsymbol{\sigma}_2)\boldsymbol{\sigma}_2^T \\ \vdots \\ P^*(\boldsymbol{\sigma}_n)\boldsymbol{\sigma}_n^T \end{bmatrix} \quad (9)$$

where $P^*(\boldsymbol{\sigma}_i) = \min(\phi, N_i \tilde{P}(\boldsymbol{\sigma}_i))$ and $N_i = 1$ if $\|\boldsymbol{\sigma}_i\| \leq \theta$ and 0 otherwise. In all of the examples below, we test the efficacy of the asymmetry correction for a range of values for θ and ϕ , including $\theta = \infty$ and $\phi = 1$, the values for which the asymmetry correction has no effect and the RF estimate $\hat{\mathbf{g}}_{cs}$ is equivalent to the RF estimate that is corrected for correlations only $\hat{\mathbf{g}}_c$.

As a first test of the efficacy of the asymmetry correction for higher-dimensional RFs, we used the LN model to simulate responses to the time-series of natural intensities (shown again in figure 2a, along with its two-point intensity distribution in figure 2b), but for these simulations the RF was chosen to have a biphasic shape that is typical of temporal RFs in early sensory systems. We repeated the simulation while increasing the dimensionality of the stimulus from $m=5$ to a more realistic value of $m=32$ (representing, for example, temporal summation of 32 frames of a spatially uniform visual stimulus).

To explicitly examine the bias in the RF estimate due to stimulus asymmetries and the efficacy of the asymmetry correction, we compared the actual RF \mathbf{g} to the RF estimated from simulated responses to the natural stimulus after correcting for correlations only $\hat{\mathbf{g}}_c$ and after correcting for both correlations and asymmetries $\hat{\mathbf{g}}_{cs}$ with optimal values for θ and ϕ . Note that for all of the simulations shown in figure 2, the best estimates of $\hat{\mathbf{g}}_{cs}$ and $\hat{\mathbf{g}}_c$ were obtained when all principal components of the stimulus were used ($\epsilon=1$ and $m^*=m$) and, thus, only results obtained with these values are shown.

The results for an RF with $m=5$ are shown in figure 2c. The color image shows the bias in $\hat{\mathbf{g}}_{cs}$ (defined as the mean squared error between \mathbf{g} and $\hat{\mathbf{g}}_{cs}$) for different values of θ and ϕ . The estimate of the RF with the lowest error (denoted by the star) was obtained when θ was set such that the stimuli with the smallest 78% of norms were included in the estimate and the maximum asymmetry correction was set to $\phi = 10^3$. For these optimal values of θ and ϕ , the RF estimate $\hat{\mathbf{g}}_{cs}$ (shown in cyan) is identical to the actual RF (shown in green). The error in the estimate increases when the RF estimate is corrected for correlations only ($\theta = \infty$ and

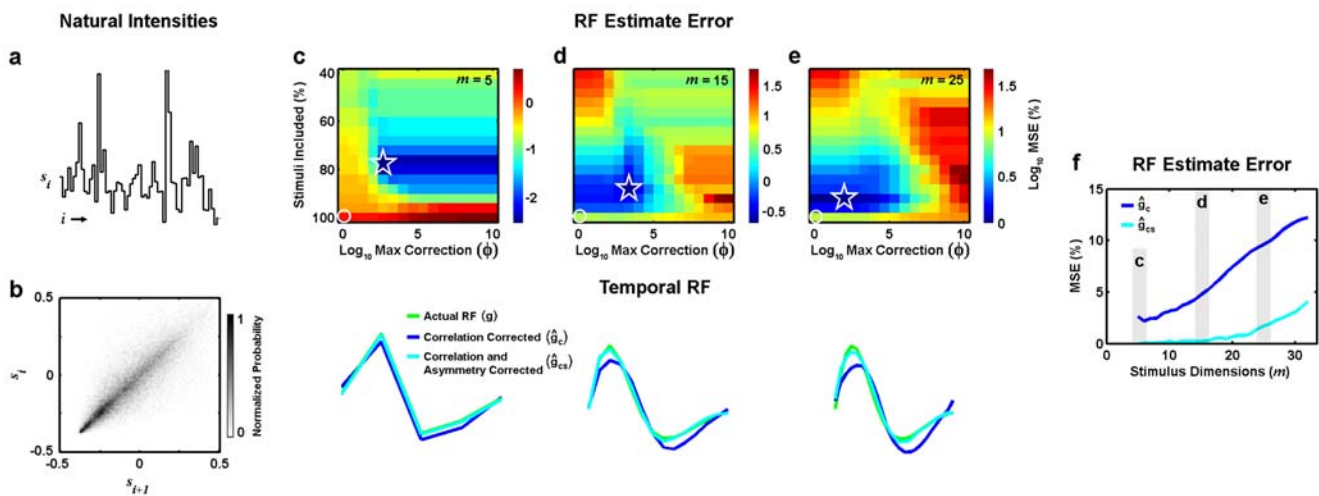


Figure 2. Temporal RF estimates from simulated responses to natural stimuli. a) 60 sample segment of a time-series of natural intensities taken from the database of van Hateren [16]. b) Two-point distribution of intensities in the natural stimulus. c) Color image showing the (\log_{10}) mean squared error between the actual RF used to simulate responses of the LN model to the natural stimulus and RF estimates corrected for stimulus correlations and asymmetries $\hat{\mathbf{g}}_{cs}$ computed from the simulated responses with different values of θ , defined in terms of the % of stimuli that were included in the RF estimate, and ϕ , the maximum value for the asymmetry correction. The minimum error is denoted by the star, and the error corresponding to correction for correlations only is denoted by the circle. The actual RF ($m=5$) used to simulate the LN model responses to the natural stimulus is shown (green), along with the RF estimated from simulated responses after correction for stimulus correlations only $\hat{\mathbf{g}}_c$ (blue) and correction for stimulus correlations and asymmetries $\hat{\mathbf{g}}_{cs}$ with optimal values of θ and ϕ (cyan). d,e) Results for $m=15$ and $m=25$, presented as in c. f) Mean squared error between actual RFs and RF estimates $\hat{\mathbf{g}}_c$ (blue) and $\hat{\mathbf{g}}_{cs}$ with optimal values of θ and ϕ (cyan) for different values of m . doi:10.1371/journal.pone.0003060.g002

$\phi = 1$, denoted by the circle) and the RF estimate $\hat{\mathbf{g}}_{\mathbf{c}}$ (shown in blue) is biased. For $m = 5$, the bias in $\hat{\mathbf{g}}_{\mathbf{c}}$ due to stimulus asymmetries is small, but this bias increases as the dimensionality of the RF increases (shown for $m = 15$ in figure 2d and $m = 25$ in figure 2e) and the asymmetry correction substantially reduces this bias. A summary of the bias in $\hat{\mathbf{g}}_{\mathbf{c}}$ and the efficacy of the asymmetry correction across a range of values for m is shown in figure 2f. The bias in $\hat{\mathbf{g}}_{\mathbf{c}}$ (blue) increases steadily as the dimensionality of the stimulus increases, while the bias in $\hat{\mathbf{g}}_{\mathbf{cs}}$ (cyan) increases much more slowly.

Spatial receptive field estimation from simulated responses to natural stimuli

When next tested the efficacy of the asymmetry correction for even higher-dimensional RFs. We used the LN model to simulate responses to a series of natural images taken from the database of van Hateren [33]. For these simulations, the RF g was defined by m points in space with a center-surround structure that is typical of spatial RFs in the early visual pathway and the stimulus \mathbf{s}_i was defined as

$$\mathbf{s}_i = \begin{bmatrix} s_{i1} \\ s_{i2} \\ \vdots \\ s_{im} \end{bmatrix}$$

where s_{ij} is the intensity of pixel j in image i . Several example images are shown in figure 3a. We again repeated the simulation while increasing the dimensionality of the RF from $m = 81$ to $m = 625$, and the natural images were resized to achieve the desired dimensionality.

The results for an RF with $m = 81$ are shown in figure 3b. In this example, the estimate of the RF with the lowest error (denoted by the star) was obtained when all principal components of the stimulus were used ($\epsilon = 1$ and $m^* = 81$) and θ was set such that the stimuli with the smallest 80% of norms were included in the estimate and the maximum asymmetry correction was set to $\phi = 10^6$. The error in the estimate increases when the RF estimate is corrected for correlations only ($\theta = \infty$ and $\phi = 1$, denoted by the circle). The RF estimate $\hat{\mathbf{g}}_{\mathbf{c}}$ that is corrected for correlations only has a weaker surround than the actual RF \mathbf{g} , while the RF estimate corrected for both correlations and asymmetries $\hat{\mathbf{g}}_{\mathbf{cs}}$ and the actual RF are similar.

Figure 3c shows the results for a high-dimensional RF with $m = 576$. In this example, the estimate of the RF with the lowest error was obtained when only the principal components necessary to explain 73% of the variance in the stimulus were used ($\epsilon = 0.73$ and $m^* = 75$; in this example, these values were optimal for both $\hat{\mathbf{g}}_{\mathbf{cs}}$ and $\hat{\mathbf{g}}_{\mathbf{c}}$), with all stimuli included in the estimate and the maximum asymmetry correction set to $\phi = 10^5$. When only correlations are corrected, the RF estimate $\hat{\mathbf{g}}_{\mathbf{c}}$ has a strong bias evidenced by the large negative values at the bottom of the RF and correcting for both correlations and asymmetries in $\hat{\mathbf{g}}_{\mathbf{cs}}$ reduces this bias. A summary of the bias in $\hat{\mathbf{g}}_{\mathbf{c}}$ and the efficacy of the asymmetry correction across a range of values m is shown in figure 3d. As with the temporal RF examples shown in figure 2, the bias in the spatial RF estimate after correcting for correlations only $\hat{\mathbf{g}}_{\mathbf{c}}$ (blue) increases steadily as the dimensionality of the stimulus increases, while the bias in the estimate after correcting for both correlations and asymmetries $\hat{\mathbf{g}}_{\mathbf{cs}}$ (cyan) increases much more slowly.

Temporal receptive field estimation from experimental responses to natural stimuli

The above results demonstrate that correcting for stimulus asymmetries can reduce the bias in an RF estimate computed from simulated responses to a natural stimulus, even for relatively high-dimensional RFs. However, under experimental conditions, the data available for RF estimation can be limited to a relatively small number of noisy observations of the neural response. To determine whether explicit correction for stimulus asymmetries is sufficient to provide accurate RF estimates under such conditions, we recorded retinal responses to the same natural time-series used in the simulated examples described in figure 2 (shown again in figure 4a). The spatially uniform stimulus was projected onto an isolated salamander retina and action potentials from ganglion cells were recorded extracellularly. The methods for these experiments have been described in detail previously [34].

The experimental responses were used to estimate RFs with correction for stimulus correlations and asymmetries as described above. Because this is an experimental situation, we do not have access to the actual RF with which to compare our estimate. Instead, we evaluate the quality of the RF estimates by measuring their ability to predict experimental responses to novel natural stimuli. First, the stimulus/response data for each neuron are divided into ‘training’ and ‘testing’ segments to avoid contamination of the evaluation from ‘overfitting’ the noise in the response. Next, the training data are used to estimate the RF and determine the optimal value of ϵ for $\hat{\mathbf{g}}_{\mathbf{c}}$, and the optimal values of ϵ , θ , and ϕ for $\hat{\mathbf{g}}_{\mathbf{cs}}$. The optimal parameter values are chosen by using the RF estimate in the LN model to simulate responses to the training stimulus (with $f(\cdot)$ also estimated from the training data as in [20]) and maximizing the predictive power measured as the correlation coefficient between the simulated and actual responses. Finally, the resulting RF estimates are used to simulate responses to the testing stimulus and the predictive power between the simulated and actual responses is measured. To provide an additional benchmark for comparison, we also recorded responses of the same cells to a spatially uniform Gaussian white noise stimulus and computed the reverse correlation RF estimate $\hat{\mathbf{g}}$ from these responses.

The results for an example OFF-center cell are shown in figures 4b and c with $m = 18$. For this cell, the estimate of the RF that had the highest predictive power for the training data was obtained when all principal components of the stimulus were used ($\epsilon = 1$ and $m^* = 18$), with θ set such that the stimuli with the smallest 77% of norms were included in the estimate and the maximum asymmetry correction set to $\phi = 10^3$. When only correlations are corrected, the predictive power of the RF estimate for the training data is decreased and differences between $\hat{\mathbf{g}}_{\mathbf{cs}}$ (cyan) and $\hat{\mathbf{g}}_{\mathbf{c}}$ (blue) are evident in the later phases of the RF estimates. Additionally, both $\hat{\mathbf{g}}_{\mathbf{cs}}$ and $\hat{\mathbf{g}}_{\mathbf{c}}$ are substantially different from the RF estimated from responses to the white noise stimulus $\hat{\mathbf{g}}$.

To evaluate the quality of the different RF estimates, we measured their predictive power for the testing data as described above. As shown in figure 4d for a sample of 10 cells, the predictive power of the RFs estimated from responses to the natural stimulus with correction for stimulus correlations and asymmetries $\hat{\mathbf{g}}_{\mathbf{cs}}$ is significantly larger than that of the RFs estimated from responses to the natural stimulus with correction for stimulus correlations only $\hat{\mathbf{g}}_{\mathbf{c}}$, as well as that of the RFs estimated from responses to the white noise stimulus $\hat{\mathbf{g}}$ (paired t-tests, $p < 0.01$). On average, the predictive power of $\hat{\mathbf{g}}_{\mathbf{cs}}$ was 5% larger than that of $\hat{\mathbf{g}}_{\mathbf{c}}$ and 11% larger than that of $\hat{\mathbf{g}}$, with increases for individual cells as large as 13% and 28%, respectively. This suggests that the RFs with correction for stimulus correlations and asymmetries do indeed provide a more accurate description of

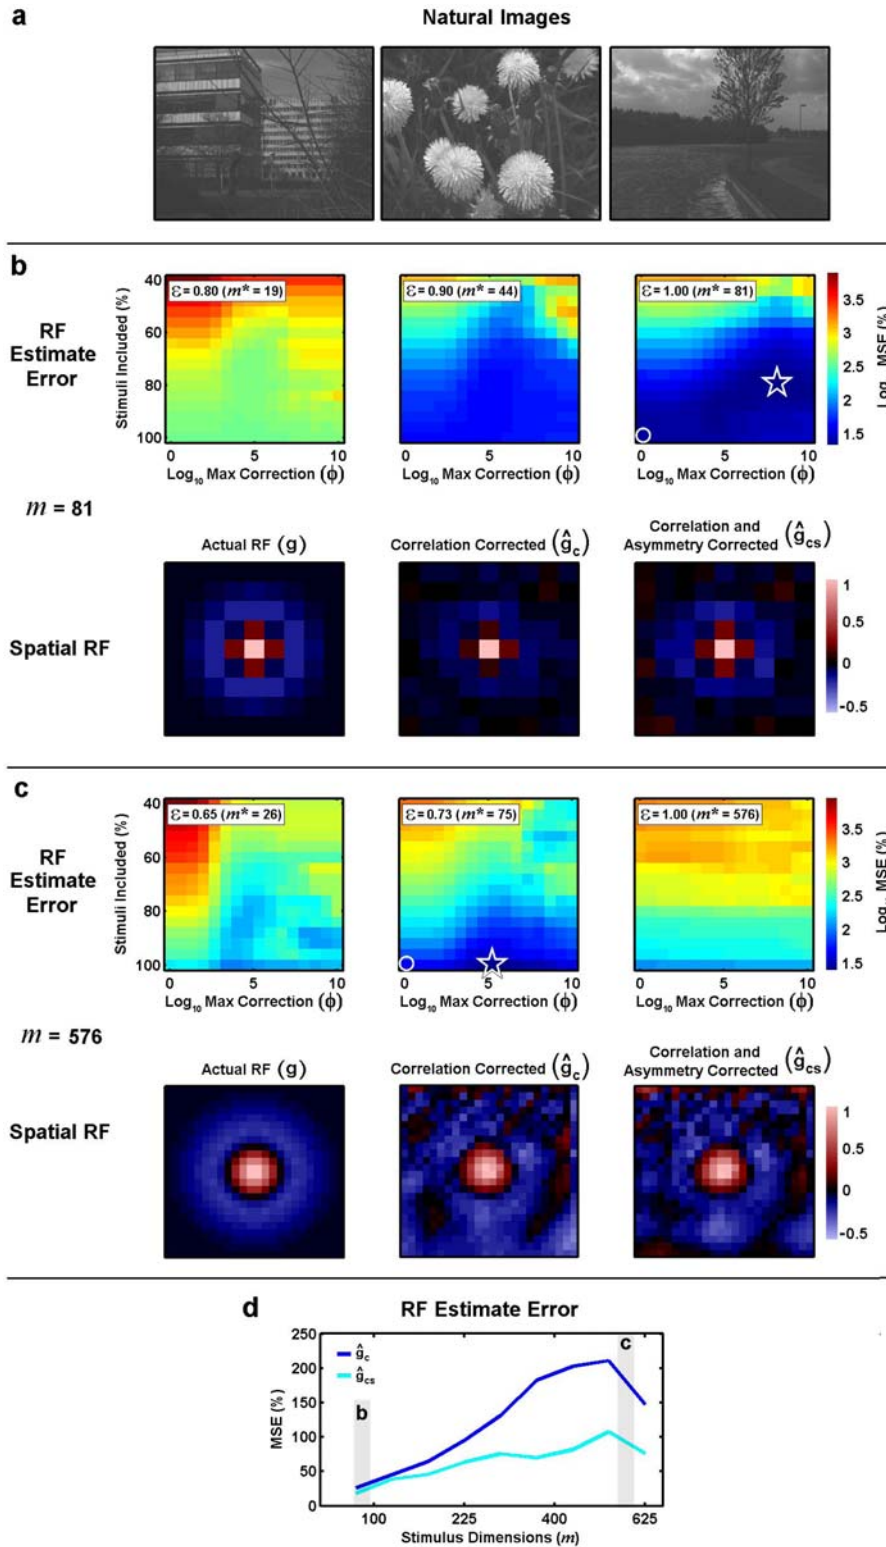


Figure 3. Spatial RF estimates from simulated responses to natural stimuli. a) Three example natural images taken from the database of van Hateren [33]. b) Color image showing the mean squared error between the actual RF used to simulate responses of the LN model to the natural stimulus and RF estimates corrected for stimulus correlations and asymmetries \hat{g}_{cs} computed from the simulated responses with different values of ϵ , the fraction of the variance explained by the principal components of the stimulus used in computing the RF estimate, θ , defined in terms of the % of stimuli that were included in the RF estimate, and ϕ , the maximum value for the asymmetry correction. The minimum error is denoted by the star, and the error corresponding to correction for correlations only is denoted by the circle. The actual RF ($m=81$) used to simulate the LN model responses to the natural stimulus is shown, along with the RF estimated from simulated responses for stimulus correlations only \hat{g}_c with an optimal value of ϵ , and correction for stimulus correlations and asymmetries \hat{g}_{cs} with optimal values of ϵ , θ , and ϕ . c) Results for $m=576$, presented as in b. d) Mean squared error between actual RFs and RF estimates \hat{g}_c (blue) and \hat{g}_{cs} with optimal values of θ and ϕ (cyan) for different values of m .

doi:10.1371/journal.pone.0003060.g003

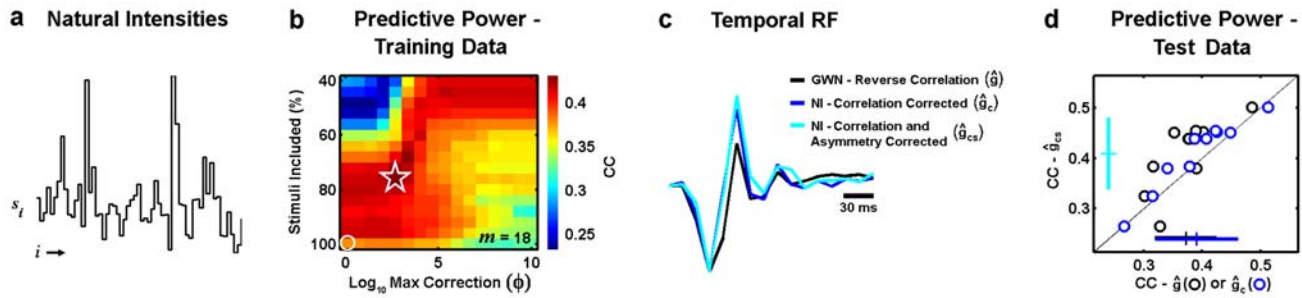


Figure 4. Temporal RF estimates from experimental responses to natural stimuli. a) 60 sample segment of a time-series of natural intensities taken from the database of van Hateren [16]. b) Color image showing the correlation coefficient (CC) between the experimental responses of a retinal ganglion cell to the ‘training’ segment of the natural stimulus and the responses predicted by RF estimates ($m = 18$) corrected for stimulus correlations and asymmetries \hat{g}_{cs} computed from the experimental responses with different values of θ , defined in terms of the % of stimuli that were included in the RF estimate, and ϕ , the maximum value for the asymmetry correction. The maximum CC is denoted by the star, and the CC corresponding to correction for correlations only is denoted by the circle. c) The RF estimated from experimental responses to the natural stimuli after correction for stimulus correlations only \hat{g}_c is shown (blue), along with the RF estimated from experimental responses to the natural stimuli after correction for stimulus correlations and asymmetries \hat{g}_{cs} with optimal values of θ and ϕ (cyan) and the RF estimated from experimental responses of the same cell to white noise stimuli using reverse correlation \hat{g} (black). d) Correlation coefficient between predicted and actual firing for responses to the ‘testing’ segment of the natural stimulus for RF estimates \hat{g}_c , \hat{g}_{cs} with optimal values of θ and ϕ , and \hat{g} for a sample of 10 retinal ganglion cells. Crosses indicate sample mean and standard deviation. doi:10.1371/journal.pone.0003060.g004

temporal processing in retinal ganglion cells than the RFs estimated with correction for stimulus correlations only and that the explicit correction for stimulus asymmetries can be effective in the analysis of experimental responses to natural stimuli.

Discussion

We have described a method for correcting the bias in reverse correlation RF estimates that arises from the asymmetries typical of natural stimuli. Using simulated neural responses, we have illustrated how stimulus asymmetries can bias reverse correlation RF estimates (even for uncorrelated stimuli) and demonstrated how explicit correction for spherical asymmetries can remove this bias. We have also shown that this method is suitable for estimating RFs under experimental conditions using retinal responses to natural stimuli. Below, we discuss the limitations of the asymmetry correction method presented here and consider other methods for RF estimation from responses to natural stimuli.

Limitations of the asymmetry correction

The primary limitation of the asymmetry correction method, at least in theory, is that its efficacy decreases as the stimulus dimensionality increases. Because the asymmetry correction method requires that responses to stimuli with low probability are weighted heavily in the RF estimate, the stability of the correction suffers as the stimulus dimensionality (and, thus, for natural stimuli, the percentage of stimuli with low probability) is increased. For high-dimensional stimuli, this problem is mitigated somewhat by using only a subset of the principal components of the stimulus in computing the transformation that corrects the RF estimate for second-order stimulus correlations. We further addressed this problem by introducing two additional parameters, θ and ϕ , into the asymmetry correction that limit the set of stimuli used in the RF estimate and the maximum value of the asymmetry correction. It is important to note that if the system is indeed well described by the LN model (the assumption upon which most RF-based analyses are based), then the exclusion of any particular subset of stimuli will not bias the RF estimate. Furthermore, as long as the range of values of θ and ϕ that are tested include $\theta = \infty$ and $\phi = 1$, then the resulting

RF estimate will be at worst equivalent to the RF estimate that is corrected for correlations only.

The asymmetry correction method presented here is only useful when the neural response can be accurately described by the standard LN model. However, there are many sensory neurons with nonlinear response properties for which the standard LN model is inadequate. Several recent studies have used spike-triggered covariance (STC) techniques to estimate RFs for more sophisticated LN models containing multiple linear filters with nonlinear interactions (for review, see [18,22,35]). These techniques have allowed for the characterization of neural responses that are incompatible with the standard LN model, such as motion sensitive cells in the fly lobula plate and complex cells in the primary visual cortex [3,4,36,37]. The correlations and asymmetries in natural stimuli can also bias estimates of the STC (indeed, the symmetry requirements for unbiased STC estimates are stricter than those for RF estimates in the standard LN model [18,21]). It is possible that the explicit asymmetry correction introduced here could be extended to provide unbiased STC estimates from responses to natural stimuli.

The standard LN model also assumes that the structure of the RF is time-invariant. While this assumption is appropriate for many sensory neurons under stationary stimulus conditions, the natural environment can be highly nonstationary and changes in the statistical properties of the stimulus can evoke adaptive changes in neural response properties (for reviews, see [38,39]). This adaptation is reflected in significant changes in the structure of the RF and, as a result, the standard LN model with a time-invariant RF is often insufficient to describe neural responses to nonstationary stimuli. To track adaptive changes in RF structure, we have previously extended the least-squares approach for RF estimation (see equation 2) to estimate time-varying RFs [30,31,40]. As described above, the least-squares approach corrects RF estimates for the bias introduced by stimulus correlations, but not for the bias introduced by stimulus asymmetries. The asymmetry correction method developed here could also be extended to estimate time-varying RFs, but the stability of the correction may depend strongly on the degree of nonstationarity in the distribution of stimulus intensities.

Other methods for receptive field estimation from responses to natural stimuli

In a number of studies in both the visual and auditory systems, RFs have been estimated from responses to natural stimuli and the bias due to second-order stimulus correlations has been corrected using a least-squares approach (equation 2, or some variant thereof) [3,4,9,11,27–31]. Detailed descriptions and analyses of this approach have also been published [21,23–25]. In the studies cited above, the biases in the RF estimates due to stimulus asymmetries were not explicitly corrected. However, in some of these studies, the authors attempted to quantify the bias in the RF estimates introduced by stimulus asymmetries by simulating responses to natural stimuli using an LN model with a known RF and comparing the actual RF to the RF estimated from the simulated responses [3,4,31,41]. Because the actual and estimated RFs were similar, the authors concluded that the bias in the RF estimates introduced by stimulus asymmetries was minimal. In contrast, a similar comparison between RF estimates corrected for correlations only and actual RFs in our results revealed large differences (see, for example, figure 2e) and another recent simulation study found similar results [26]. Taken together, these results imply that the bias in RF estimates introduced by stimulus asymmetries is dependent on the specific statistical properties of the stimulus and suggest that these effects should be investigated explicitly in each new experimental context.

Recently, several new techniques for RF estimation have been developed that use gradient descent methods to produce RF estimates that are independent of both stimulus correlations and asymmetries. One set of techniques minimizes the same cost function used in reverse correlation RF estimates, the mean squared error between the predicted and actual responses [32,42], while another set maximizes some variant of the mutual information between the stimulus and the predicted response

[21,43,44]. Both sets of techniques have already been used successfully to estimate high-dimensional RFs under experimental conditions. The only potential drawback to these approaches is that they require a search algorithm to determine the optimal RF and it may be difficult to avoid local optima. Another promising new approach involves maximum likelihood estimation of a parametric LN model [45]. While this approach also requires a search algorithm to find the optimal parameters, it takes advantage of the fact that, for certain forms of the LN model, the likelihood surface is convex, allowing for an efficient search. This approach has been used successfully to characterize retinal ganglion cell responses to white noise stimuli [46], but has not yet been tested with natural stimuli.

While not explicitly biased by asymmetries in the stimulus, RFs estimated using gradient descent methods will, by definition, be most successful in minimizing or maximizing the relevant cost function for predicted responses to those stimuli that occur most frequently. Thus, these estimates are still influenced by the probability distribution of the stimulus and it is possible that they could also benefit from a correction similar to the asymmetry correction described here in which the contribution of each response to the cost function is weighted by the probability of the corresponding stimulus.

Acknowledgments

The authors would like to thank Christian Leibold and Hamish Meffin for helpful conversations.

Author Contributions

Conceived and designed the experiments: NAL GBS TH. Performed the experiments: NAL TI. Analyzed the data: NAL. Wrote the paper: NAL TH.

References

- Dan Y, Atick JJ, Reid RC (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J Neurosci* 16: 3351–3362.
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287: 1273–1276.
- Touryan J, Felsen G, Dan Y (2005) Spatial structure of complex cell receptive fields measured with natural images. *Neuron* 45: 781–791.
- Felsen G, Touryan J, Han F, Dan Y (2005) Cortical sensitivity to visual features in natural scenes. *PLoS Biol* 3: e342.
- Rieke F, Bodnar DA, Bialek W (1995) Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents. *Proc Biol Sci* 262: 259–265.
- Escabi MA, Miller LM, Read HL, Schreiner CE (2003) Naturalistic auditory contrast improves spectrotemporal coding in the cat inferior colliculus. *J Neurosci* 23: 11489–11504.
- Hsu A, Woolley SM, Fremouw TE, Theunissen FE (2004) Modulation power and phase spectrum of natural sounds enhance neural encoding performed by single auditory neurons. *J Neurosci* 24: 9201–9211.
- Garcia-Lazaro JA, Ahmed B, Schnupp JW (2006) Tuning to natural stimulus dynamics in primary auditory cortex. *Curr Biol* 16: 264–271.
- Theunissen FE, Sen K, Doupe AJ (2000) Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J Neurosci* 20: 2315–2331.
- Klug A, Bauer EE, Hanson JT, Hurley L, Meitzen J, Pollak GD (2002) Response selectivity for species-specific calls in the inferior colliculus of Mexican free-tailed bats is generated by inhibition. *J Neurophysiol* 88: 1941–1954.
- David SV, Vinje WE, Gallant JL (2004) Natural stimulus statistics alter the receptive field structure of v1 neurons. *J Neurosci* 24: 6991–7006.
- Field DJ (1987) Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A* 4: 2379–2394.
- Ruderman DL, Bialek W (1994) Statistics of natural images: Scaling in the woods. *Phys Rev Lett* 73: 814–817.
- Dong DW, Atick JJ (1995) Statistics of Natural Time-Varying Images. *Network-Computation in Neural Systems* 6: 345–358.
- Singh NC, Theunissen FE (2003) Modulation spectra of natural sounds and ethological theories of auditory processing. *J Acoust Soc Am* 114: 3394–3411.
- van Hateren JH (1997) Processing of natural time series of intensities by the visual system of the blowfly. *Vision Res* 37: 3407–3416.
- Thomson MG (1999) Visual coding and the phase structure of natural scenes. *Network* 10: 123–132.
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vis* 6: 484–507.
- Bussgang JJ (1952) Crosscorrelation functions of amplitude distorted Gaussian signals. *MIT Res Lab Elec Tech Rep*. pp 1–14.
- Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. *Network* 12: 199–213.
- Paninski L (2003) Convergence properties of three spike-triggered analysis techniques. *Network* 14: 437–464.
- Simoncelli EP, Paninski L, Pillow J, Schwartz O (2003) Characterization of neural responses with stochastic stimuli. In: *The Cognitive Neurosciences*. MIT Press. pp. 23.
- Theunissen FE, David SV, Singh NC, Hsu A, Vinje WE, Gallant JL (2001) Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12: 289–316.
- Willmore B, Smyth D (2003) Methods for first-order kernel estimation: simple-cell receptive fields from responses to natural scenes. *Network* 14: 553–577.
- David SV, Gallant JL (2005) Predicting neuronal responses during natural vision. *Network* 16: 239–260.
- Christianson GB, Sahani M, Linden JF (2008) The consequences of response nonlinearities for interpretation of spectrotemporal receptive fields. *J Neurosci* 28: 446–455.
- Ringach DL, Hawken MJ, Shapley R (2002) Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *J Vis* 2: 12–24.
- Machens CK, Wehr MS, Zador AM (2004) Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24: 1089–1100.
- Smyth D, Willmore B, Baker GE, Thompson ID, Tolhurst DJ (2003) The receptive-field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *J Neurosci* 23: 4746–4759.
- Lesica NA, Stanley GB (2006) Decoupling functional mechanisms of adaptive encoding. *Network* 17: 43–60.
- Lesica NA, Jin J, Weng C, Yeh CI, Butts DA, Stanley GB, Alonso JM (2007) Adaptation to Stimulus Contrast and Correlations during Natural Visual Stimulation. *Neuron* 55: 479–491.

32. David SV, Mesgarani N, Shamma SA (2007) Estimating sparse spectro-temporal receptive fields with natural stimuli. *Network* 18: 191–212.
33. van Hateren JH, Ruderman DL (1998) Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc Biol Sci* 265: 2315–2320.
34. Hosoya T, Baccus SA, Meister M (2005) Dynamic predictive coding by the retina. *Nature* 436: 71–77.
35. Ringach DL (2004) Mapping receptive fields in primary visual cortex. *J Physiol* 558: 717–728.
36. Brenner N, Bialek W, de Ruyter van SR (2000) Adaptive rescaling maximizes information transmission. *Neuron* 26: 695–702.
37. Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2005) Spatiotemporal elements of macaque v1 receptive fields. *Neuron* 46: 945–956.
38. Meister M, Berry MJ (1999) The neural code of the retina. *Neuron* 22: 435–450.
39. Clifford CW, Webster MA, Stanley GB, Stocker AA, Kohn A, Sharpee TO, Schwartz O (2007) Visual adaptation: Neural, psychological and computational aspects. *Vision Res* 47: 3125–3131.
40. Lesica NA, Stanley GB (2005) Improved tracking of time-varying encoding properties of visual neurons by extended recursive least-squares. *IEEE Trans Neural Syst Rehabil Eng* 13: 194–200.
41. Lesica NA, Grothe B (2008) Efficient temporal processing of naturalistic sounds. *PLoS ONE* 3: e1655.
42. Sahani M, Linden JF (2003) Evidence Optimization Techniques for Estimating Stimulus-Response Functions. In: Becker S, Thrun S, Obermayer K, eds. *Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press. pp 317–324.
43. Sharpee T, Rust NC, Bialek W (2004) Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* 16: 223–250.
44. Sharpee TO, Sugihara H, Kurgansky AV, Rebrik SP, Stryker MP, Miller KD (2006) Adaptive filtering enhances information transmission in visual cortex. *Nature* 439: 936–942.
45. Paninski L (2004) Maximum likelihood estimation of cascade point-process neural encoding models. *Network* 15: 243–262.
46. Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ (2005) Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *J Neurosci* 25: 11003–11013.