

RESEARCH ARTICLE

# Inferring hidden structure in multilayered neural circuits


Niru Maheswaranathan<sup>1aa</sup>, David B. Kastner<sup>1ab</sup>, Stephen A. Baccus<sup>2</sup>, Surya Ganguli<sup>3\*</sup>

**1** Neurosciences Graduate Program, Stanford University, Stanford, California, United States of America, **2** Department of Neurobiology, Stanford University, Stanford, California, United States of America, **3** Department of Applied Physics, Stanford University, Stanford, California, United States of America

<sup>aa</sup> Current address: Google Brain, Mountain View, California, United States of America

<sup>ab</sup> Current address: Department of Psychiatry, UCSF, San Francisco, California, United States of America

\* [sganguli@stanford.edu](mailto:sganguli@stanford.edu)



## Abstract

A central challenge in sensory neuroscience involves understanding how neural circuits shape computations across cascaded cell layers. Here we attempt to reconstruct the response properties of experimentally unobserved neurons in the interior of a multilayered neural circuit, using cascaded linear-nonlinear (LN-LN) models. We combine non-smooth regularization with proximal consensus algorithms to overcome difficulties in fitting such models that arise from the high dimensionality of their parameter space. We apply this framework to retinal ganglion cell processing, learning LN-LN models of retinal circuitry consisting of thousands of parameters, using 40 minutes of responses to white noise. Our models demonstrate a 53% improvement in predicting ganglion cell spikes over classical linear-nonlinear (LN) models. Internal nonlinear subunits of the model match properties of retinal bipolar cells in both receptive field structure and number. Subunits have consistently high thresholds, suppressing all but a small fraction of inputs, leading to sparse activity patterns in which only one subunit drives ganglion cell spiking at any time. From the model's parameters, we predict that the removal of visual redundancies through stimulus decorrelation across space, a central tenet of efficient coding theory, originates primarily from bipolar cell synapses. Furthermore, the composite nonlinear computation performed by retinal circuitry corresponds to a boolean OR function applied to bipolar cell feature detectors. Our methods are statistically and computationally efficient, enabling us to rapidly learn hierarchical non-linear models as well as efficiently compute widely used descriptive statistics such as the spike triggered average (STA) and covariance (STC) for high dimensional stimuli. This general computational framework may aid in extracting principles of nonlinear hierarchical sensory processing across diverse modalities from limited data.

## OPEN ACCESS

**Citation:** Maheswaranathan N, Kastner DB, Baccus SA, Ganguli S (2018) Inferring hidden structure in multilayered neural circuits. *PLoS Comput Biol* 14(8): e1006291. <https://doi.org/10.1371/journal.pcbi.1006291>

**Editor:** Peter E. Latham, UCL, UNITED KINGDOM

**Received:** March 20, 2017

**Accepted:** June 9, 2018

**Published:** August 23, 2018

**Copyright:** © 2018 Maheswaranathan et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All code and data used in this paper are available at: <https://github.com/baccuslab/inferring-hidden-structure-retinal-circuits>.

**Funding:** Stephen A. Baccus received funding from National Institutes of Health (NIH) R01EY025087, R01EY022933. Surya Ganguli received funding from the Burroughs-Wellcome, Simons, McKnight, and James S. McDonnell Foundations. Stephen A. Baccus and Surya Ganguli received a grant from Stanford University's Cracking the Neural Code program. Niru Maheswaranathan was supported by an NSF Graduate Research Fellowship (GRFP)

## Author summary

Computation in neural circuits arises from the cascaded processing of inputs through multiple cell layers. Each of these cell layers performs operations such as filtering and thresholding in order to shape a circuit's output. It remains a challenge to describe both

and the Stanford Mind, Brain, and Computation program. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

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

the computations and the mechanisms that mediate them given limited data recorded from a neural circuit. A standard approach to describing circuit computation involves building quantitative encoding models that predict the circuit response given its input, but these often fail to map in an interpretable way onto mechanisms within the circuit. In this work, we build two layer linear-nonlinear cascade models (LN-LN) in order to describe how the retinal output is shaped by nonlinear mechanisms in the inner retina. We find that these LN-LN models, fit to ganglion cell recordings alone, identify filters and nonlinearities that are readily mapped onto individual circuit components inside the retina, namely bipolar cells and the bipolar-to-ganglion cell synaptic threshold. This work demonstrates how combining simple prior knowledge of circuit properties with partial experimental recordings of a neural circuit's output can yield interpretable models of the entire circuit computation, including parts of the circuit that are hidden or not directly observed in neural recordings.

## Introduction

### Motivation

Computational models of neural responses to sensory stimuli have played a central role in addressing fundamental questions about the nervous system, including how sensory stimuli are encoded and represented, the mechanisms that generate such a neural code, and the theoretical principles governing both the sensory code and underlying mechanisms. These models often begin with a statistical description of the stimuli that precede a neural response such as the spike-triggered average (STA) [1, 2] or covariance (STC) [3–8]. These statistical measures characterize to some extent the set of effective stimuli that drive a response, but do not necessarily reveal how these statistical properties relate to cellular mechanisms or neural pathways. Going beyond descriptive statistics, an explicit representation of the neural code can be obtained by building a model to predict neural responses to sensory stimuli.

A classic approach involves a single stage of spatiotemporal filtering and a time-independent or static nonlinearity; these models include linear-nonlinear (LN) models with single or multiple pathways [1, 9–11] or generalized linear models (GLMs) with spike history feedback [12, 13]. However, these models do not directly map onto circuit anatomy and function. As a result, the interpretation of such phenomenological models, as well as how they precisely relate to underlying cellular mechanisms, remains unclear. Ideally, one would like to generate more biologically interpretable models of sensory circuits, in which sub-components of the model map in a one-to-one fashion onto cellular components of neurobiological circuits [14]. For example, model components such as spatiotemporal filtering, thresholding, and summation are readily mapped onto photoreceptor or membrane voltage dynamics, synaptic and spiking thresholds, and dendritic pooling, respectively.

A critical aspect of sensory circuits is that they operate in a hierarchical fashion in which sensory signals propagate through multiple nonlinear cell layers [15–17]. Fitting models that capture this widespread structure using neural data recorded from one layer of a circuit in response to controlled stimuli raises significant statistical and computational challenges [18–22]. A key issue is the high dimensionality of both stimulus and parameter space, as well as the existence of hidden, unobserved neurons in intermediate cell layers. The high dimensionality of parameter space can necessitate prohibitively large amounts of data and computational time required to accurately fit the model. One approach to address these difficulties is to incorporate prior knowledge about the structure and components of circuits to

constrain the model [11, 21, 23–25]. Although prior knowledge of the exact network architecture and sequence of nonlinear transformations would greatly constrain the number of possible circuit solutions, such prior knowledge is typically minimal for most neural circuits.

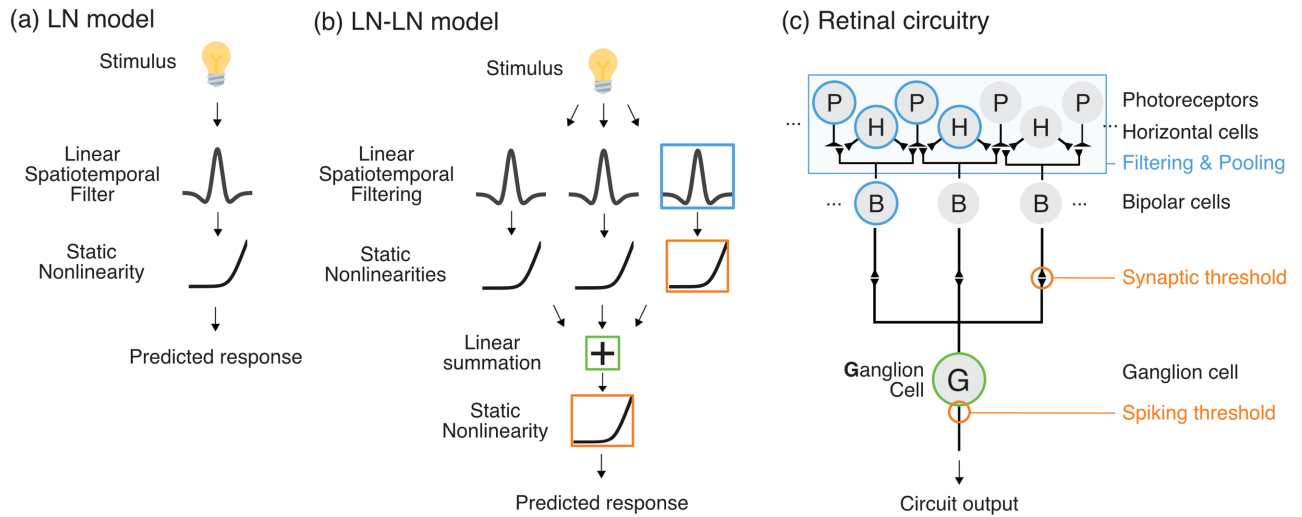
In this work, we learn hierarchical nonlinear models from recordings of ganglion cells in the salamander retina, with the goal of building more interpretable models. In particular, we focus on models with two stages of linear-nonlinear processing (LN-LN models), analogous to the specific cell layers the retina. LN-LN models have been previously proposed to describe cascaded nonlinear computation in sensory pathways such as in V1 [11, 21] and retina [18, 22, 23, 26, 27] (we elaborate on differences across these studies and our work below). In the retina, it has been proposed that there is a nonlinear transformation between bipolar and ganglion cells, however, building models that capture these nonlinearities has been a challenge due to the issues described above. Here, we find that with appropriate regularization, we are able to learn LN-LN models from recordings of ganglion cells alone that are both more accurate and more interpretable than their LN counterparts. In particular, inferred LN-LN model subunits quantitatively match properties of bipolar cells in the retina. Moreover, although the focus of this paper is on LN-LN models, we demonstrate that the algorithms we use to fit them are also useful in directly learning the STA and STC eigenvectors using very little data.

Further analysis of our learned LN-LN models reveals novel insights into retinal function, namely that: transmission between every subunit and ganglion cell pair is well described by a high threshold expansive nonlinearity (suppressing all but a small fraction of inputs), bipolar cell terminals are sparsely active, visual inputs are most decorrelated at the subunit layer, pre-synaptic to ganglion cells, and finally the composite computation performed by the retinal ganglion cell output corresponds to a boolean OR function of bipolar cell feature detectors. Collectively, these results shed light on the nature of hierarchical nonlinear computation in the retina. Our computational framework is general, however, and we hope it will aid in providing insights into hierarchical nonlinear computations across the nervous system.

## Background

The retina is a classic system for exploring the relationship between quantitative encoding models and measurements of neurobiological circuit properties [28, 29]. Signals in the retina flow from photoreceptors through populations of horizontal, bipolar, and amacrine cells before reaching the ganglion cell layer.

To characterize this complex multilayered circuitry, many studies utilize descriptive statistics such as the spike-triggered average, interpreted as the average feature encoded by a ganglion cell [1–3]. Responses are often then modeled using a linear-nonlinear (LN) framework (schematized in Fig 1A). A major reason for the widespread adoption of LN models is their high level of tractability; learning their parameters can be accomplished by solving a simple convex optimization problem [2], or alternatively, estimated using straightforward reverse correlation analyses [1]. However, LN models have two major drawbacks: it is difficult to map them onto biophysical mechanisms in retinal circuitry, and they do not accurately describe ganglion cell responses across diverse stimuli. Regarding mechanisms, the spatiotemporal linear filter of the LN model is typically interpreted as mapping onto the aggregate sequential mechanisms of phototransduction, signal filtering and transmission through bipolar and amacrine cell pathways, and summation at the ganglion cell, while the nonlinearity is mapped onto the spiking threshold of ganglion cells. Regarding accuracy, while previous studies have found that these simple models can, for some neurons, capture most of the variance of the responses to low-resolution spatiotemporal white noise [9, 12, 20], they do not describe responses to stimuli with more structure such as natural scenes [13, 30–33].



**Fig 1. Schematic of the LN-LN model and corresponding retinal circuitry.** (a) The LN-LN cascade contains a bank of linear-nonlinear (LN) subunits, whose outputs are pooled at a second linear stage before being passed through a final nonlinearity. (b,c) The LN-LN model mapped on to a retinal circuit. The first LN stage consists of bipolar cell subunits and the bipolar-to-ganglion cell synaptic threshold. The second LN stage is pooling at the ganglion cell, plus a spiking threshold. The contribution of inhibitory amacrine cells is omitted here.

<https://doi.org/10.1371/journal.pcbi.1006291.g001>

A likely reason for these drawbacks are the nonlinearities within the retina. There can be strong rectification of signals that occurs pre-synaptic to ganglion cells [15, 34–36], breaking the assumption of composite linearity in the pathway from photoreceptors just up to the ganglion cell spiking threshold [17]. Indeed, nonlinear spatial integration within ganglion cell receptive fields was first described in the cat retina [37] in Y-type ganglion cells. A hypothetical model for this computation was proposed as a cascade of two layers of linear-nonlinear operations (LN-LN) [26, 27]. If one keeps the mean luminance constant, avoiding light adaptation in photoreceptors, the first major nonlinearity is thought to lie at the presynaptic terminal of the bipolar to ganglion cell synapse. Ganglion cells pool over multiple bipolar cell inputs, each of which can be approximated as linear-nonlinear components, termed subunits of the ganglion cell. Due to the roughly linear integration [9] that occurs at bipolar cells, we (computationally) distill mechanisms in photoreceptors and inhibitory horizontal cells into a single spatiotemporal filter with positive and negative elements that gives rise to bipolar cell signals. The second LN layer corresponds to summation or pooling across multiple subunits at the ganglion cell soma, followed by a spiking threshold. The subunit nonlinearities in these models have been shown to underlie many retinal computations including latency encoding [29], object motion sensitivity [38], and sensitivity to fine spatial structure (such as edges) in natural scenes [35]. Fig 1 shows a schematic of the LN-LN cascade and its mapping onto retinal anatomy. Functionally, these models with multiple nonlinear pathways both provide a more accurate description of ganglion cell responses and are more amenable to interpretation.

Early work on characterizing these multiple pathways motivated the use of the significant eigenvectors of the spike-triggered covariance (STC) matrix as the set of features that drives a cell, focusing on low-dimensional full field flicker stimuli [10, 39] to reduce the amount of data required for accurately estimating these eigenvectors. Significant STC eigenvectors will span the same linear subspace as the true biological filters that make up the pathways feeding onto a ganglion cell [3, 40–42]. However, the precise relationship between these eigenvectors (which obey a biologically implausible orthogonality constraint) and the individual spatiotemporal filtering properties intrinsic to multiple parallel pathways in a neural circuit remains unclear.

Instead, we take the approach of directly fitting a hierarchical, nonlinear, neural model, enabling us to jointly learn a set of non-orthogonal, biophysically plausible set of pathway filters, as well as an arbitrary, flexible nonlinearity for each pathway. Much recent and complementary work on fitting such models make simplifying assumptions in order to make model fitting tractable. For example, assuming the subunits are shifted copies of a template results in models with a single convolutional subunit filter [21, 23, 24]. However, this obscures individual variability in the spatiotemporal filters of subunits of the same type across visual space, which has been shown to be functionally important in increasing retinal resolution [43]. Another common assumption is that the subunit nonlinearities have a particular form, such as quadratic [11, 25] or sigmoidal [44]. Fitting multi-layered models with convolutional filters and fixed nonlinearities has also been successfully used to describe retinal responses to natural scenes [32, 33], although this work maximizes predictive accuracy at the expense of a one-to-one mapping of model components onto retinal circuit elements. Finally, other work focuses on particular ganglion cell types with a small number of inputs [22], constrains the input stimulus to a low-dimensional subspace (such as two halves of the receptive field [45]), or constrains the coefficients of receptive fields to be non-negative [46], thus discarding known properties of the inhibitory surround. Our approach is most similar to MacFarland et. al. [18], who formulate LN-LN models for describing nonlinearities in sensory pathways. Their model formulation uses spatiotemporal filters and smooth parameterized nonlinearities with an additional sparsity regularization penalty on the filters (encouraging filters to contain few non-zero elements), fit using gradient descent. They demonstrate that these methods recover the parameters of a simulated model cell with two subunit pathways. Our work differs technically in the types of regularization penalties we apply and our use of high-dimensional spatiotemporal stimuli, and scientifically in our focus on gaining insight into the nonlinear computations underlying spatiotemporal processing in the retina.

In this work, we do not make assumptions about or place restrictions on the number or tilting of subunit filters, the shapes of the subunit nonlinearities, the sign of receptive field elements, or the stimulus dimensionality. We additionally use a low-rank regularization penalty, that encourages approximately spatiotemporally separable filters [47, 48], a property that is common to receptive fields in a wide variety of sensory systems. In order to fit these models, we use methods based on proximal consensus algorithms (described in [Methods](#)). These allow us to use this prior knowledge about model parameters to not only fit hierarchical nonlinear models, but also perform spike-triggered analyses using much less data than otherwise required.

## Results

### Learning hierarchical nonlinear models of the retinal response

Our LN-LN model architecture (schematized in [Fig 1](#)) follows previous work [18]. The stimulus is first passed through a set of LN subunits. Each subunit filter is a spatiotemporal stimulus filter, constrained to have unit norm. The subunit nonlinearity is parameterized using a set of basis functions (Gaussian bumps) that tile the input space [12, 20] (see [Methods](#)). This parameterization is flexible enough that we could learn, for each individual subunit, any smooth nonlinearity that can be expressed as a linear combination of our basis functions. The second LN layer pools subunits through weighted summation, followed by a spiking nonlinearity that we model using a parameterized soft rectifying function  $r(x) = g \log(1 + e^{x-\theta})$ . Here  $g$  is an overall gain, and  $\theta$  is a threshold. The full set of parameters for the model consists of the spatiotemporal subunit filters, the subunit nonlinearity parameters, and the gain and threshold of the final nonlinearity.

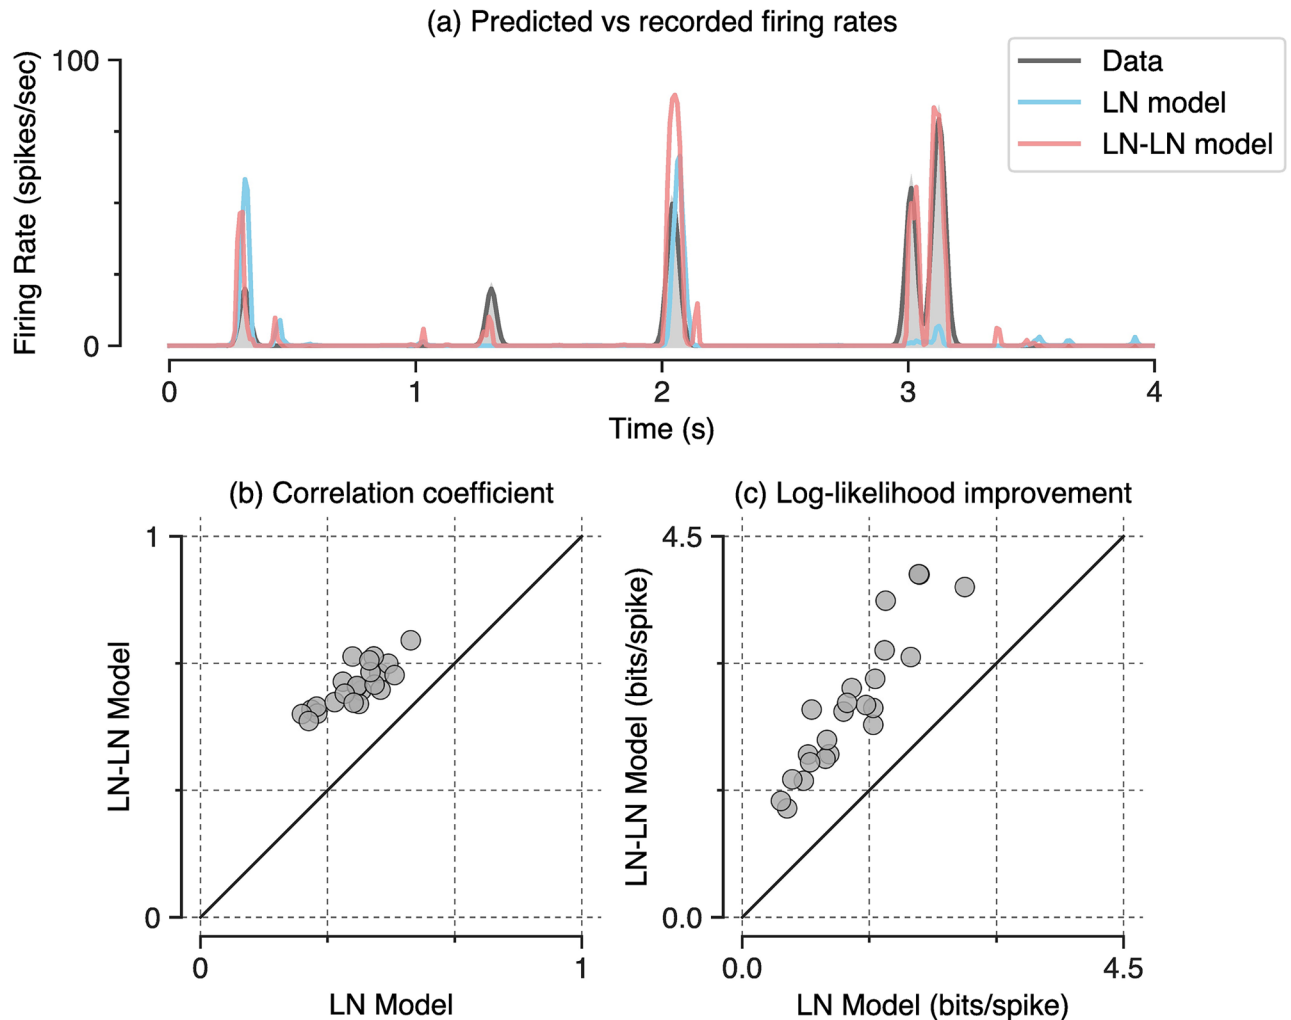
**Model fitting and performance.** We recorded ganglion cell responses to a 40 minute stimulus consisting of a 1-D spatiotemporal white noise bars stimulus. For both LN and LN-LN models we fit model parameters by optimizing the sum of the log-likelihood of recorded spikes under a Poisson noise model [12] with additional regularization terms. We found that some form of regularization was critical to prevent over-fitting of the LN-LN model to the subset of data used for training. In particular, the regularization penalties we used were  $\ell_1$  and nuclear norm penalties applied to the spatiotemporal subunit filters. The  $\ell_1$  norm encourages filters to be sparse (have few non-zero coefficients) and the nuclear norm penalty encourages the space-time filter to be low rank (comprised of a small number of features separable in space and time). We chose the weights of the  $\ell_1$  and nuclear norm regularization penalties, both for the LN and LN-LN models, through cross-validation on a small subset of cells, and then held these weights constant across all cells. Our subsequent results indicate that we do not have to fine tune these hyperparameters on a cell by cell basis to achieve good predictive performance. Finally, because different cells may have different numbers of functional subunits, for the LN-LN models, we chose the optimal number of subunits on a cell-by-cell basis by maximizing performance on held-out data through cross-validation. No additional structure was imposed on the subunits such as spatial repetition, overlap, or non-negativity. We optimize the log-likelihood and regularization terms using proximal algorithms (see [Methods](#) for an overview).

We find that the LN-LN model significantly outperforms the LN model at describing responses of ganglion cells, for all recorded cells. [Fig 2A](#) shows firing rate traces for an example cell, comparing the recorded response (gray) with an LN model (blue) and an LN-LN model (red). We quantify the similarity between predicted and recorded firing rate traces using either the Pearson correlation coefficient or the log-likelihood of held-out data under the model. All log-likelihood values are reported as an increase over the log-likelihood of a fixed mean firing rate model, scaled by the firing rate (yielding units of bits/spike). Summarized across  $n = 23$  recorded ganglion cells, we find that the LN-LN model yields a consistent improvement over the LN model using either metric ([Fig 2B and 2C](#)). Overall, this demonstrated performance improvement indicates that nonlinear spatial integration is fundamental in driving ganglion cell responses, even to white noise stimuli, and that an LN model is not sufficient to capture the response to spatiotemporal white noise [22, 49]. This salient, intermediate rectification that we identify computationally is consistent with previous measurements of bipolar-to-ganglion cell transmission in the retina [45, 50].

**Internal structure of learned models.** Given the improved performance of our hierarchical nonlinear subunit models, we examined the internal structure of the model for insights into retinal structure and computation. [Fig 3](#) shows a visualization of the parameters learned for an example cell. [Fig 3A and 3B](#) shows the parameters for the classical LN model, for comparison, while [Fig 3C and 3D](#) shows the corresponding subunit filters and subunit nonlinearities in the first stage of the LN-LN model, fit to the same cell. The subunit filters had a similar temporal structure, but smaller spatial profiles compared to that of the LN model and the nonlinearities associated with each subunit were roughly monotonic with high thresholds (quantified below). These qualitative properties were consistent across all cells.

## Physiological properties of learned LN-LN models


Here we examine, in more detail, quantitative properties of learned LN-LN models that can be compared to physiological properties of the retina. We find that model subunits quantitatively resemble bipolar cells in terms of receptive field properties and number, and that these intermediate subunits consistently have high-threshold nonlinearities.



**Fig 2. LN-LN models predict held-out response data better than LN models.** (a) Firing rates for an example neuron. The recorded firing rate (shaded, gray), is shown along with the LN model prediction (dashed, green) and the LN-LN prediction (solid, red). (b) LN-LN performance on held out data vs. the LN model, measured using correlation coefficient between the model and held out data. Note that all cells are above the diagonal. (c) Same as in (b), but with the performance metric of log-likelihood improvement over the mean rate in bits per spike.

<https://doi.org/10.1371/journal.pcbi.1006291.g002>

**Inferred hidden units quantitatively resemble bipolar cell receptive fields.** Mapping the LN-LN model onto retinal anatomy leads us to believe that the first layer filters (some examples of which are shown in Fig 3C) should mimic or capture filtering properties pre-synaptic to bipolar cells in the inner retina. To examine this possibility, we compared the first layer model filters to properties of bipolar cell receptive fields. An example learned model sub-unit receptive field is shown in Fig 4A, while a bipolar cell receptive field, obtained from direct intracellular recording of a bipolar cell, is shown in Fig 4B. Qualitatively, we found that the filters in the LN-LN model matched these bipolar cell RFs, as well as previously reported bipolar cell receptive fields [50, 51]: both had center-surround receptive fields with similar spatial extents. We further quantified the degree of space-time separability of the filters using the numerical or stable rank [52], which is a measure of rank insensitive to small amounts of noise in the matrix (a stable rank of one indicates the filter is exactly space-time separable). We found that the degree of space-time separability of recorded bipolar cell receptive fields and




**Fig 3. Example LN-LN model parameters fit to a recording of an OFF retinal ganglion cell.** (a and b): LN-model parameters, consisting of a single spatial filter (a) and nonlinearity (b). (c and d) LN-LN model parameters. (c) First layer filters (top) and nonlinearities (bottom) of an LN-LN model fit to the same cell. Spatial profiles of filters are shown in gray to the right of the filters. The subunit filters have a much smaller spatial extent compared to the LN filter, but similar temporal profiles.

<https://doi.org/10.1371/journal.pcbi.1006291.g003>

inferred model subunits were quite similar ( $1.28 \pm 0.01$  and  $1.39 \pm 0.03$ , respectively), indicating that the nuclear norm penalty was not artificially reducing the rank of our model filters. This also demonstrates the advantage of using a soft rank penalty, such as the nuclear norm, as opposed to explicitly constraining the rank to any specific integer.

To quantitatively compare these model-derived and ground-truth bipolar cell receptive fields, we fit the spatial receptive field with a difference of Gaussians function to estimate the RF center and surround sizes. We find the RF centers for the LN-LN subunit filters are much smaller than the corresponding LN model filter. Furthermore, the size of these LN-LN subunit



**Fig 4. Comparison of subunit filter parameters with intracellular bipolar cell recordings.** (a) An example subunit bipolar cell. (b) A recorded bipolar cell receptive field. (c) Receptive field centers sizes for subunit filters (blue), LN model filters (green), and recorded bipolar cells (black point). (d) Same as in (b), but with receptive field surround sizes.

<https://doi.org/10.1371/journal.pcbi.1006291.g004>




centers matched the size of the RF center measured from intracellular recordings of  $n = 8$  bipolar cells (Fig 4C). The recorded bipolar cells, LN model filters (ganglion cells), and LN-LN subunit filters all had similar surround sizes (Fig 4D). More example bipolar cell receptive fields are provided in S1 Fig.

Note that this match between LN-LN model subunits and the RF properties of bipolar cells was not a pre-specified constraint placed on our model, but instead arose as an emergent property of predicting ganglion cell responses to white noise stimuli. These results indicate that our modeling framework not only enables higher performing predictive models of the retinal response, but can also reconstruct important aspects of the unobserved interior of the retina.

**Number of inferred subunits.** The number of subunits utilized in the LN-LN model for any individual cell was chosen to optimize model predictive performance on a held-out data set via cross-validation. That is, we fit models with different numbers of subunits and selected the one with the best performance on a validation set. We find that for models with more subunits than necessary, extra subunits are ignored (the learned nonlinearity for these subunits is flat, thus they do not modulate the firing rate).

Fig 5A shows the model performance, quantified as the difference between the LN-LN model and the LN model, across a population of cells as a function of the number of subunits included in the LN-LN model. We find that models with four to six subunits maximized model performance on held-out data. Note that the stimuli used here are one-dimensional spatiotemporal bars that have constant luminance across one spatial dimension. Thus each model subunit likely corresponds to the combination of multiple bipolar cell inputs whose receptive fields overlap a particular bar in the stimulus. Previous anatomical studies of bipolar cell density and axonal branching width [53, 54] as well as functional studies [46] suggest that a typical ganglion cell in the salamander retina receives input from 10–50 bipolar cells whose receptive fields are tiled across two dimensional space. The number of independently activated groups of such a two dimensional array of bipolar cells, in response to a one dimensional bar stimulus is then expected to be reduced from the total number of bipolar cells, roughly, by a square root



**Fig 5. LN-LN model parameter analysis.** (a) Performance improvement (increase in correlation coefficient relative to an LN model) as a function of the number of subunits used in the LN-LN model. Error bars indicate the standard error across 23 cells. (b) Subunit nonlinearities learned across all ganglion cells. For reference the white noise input to a subunit nonlinearity has standard deviation 1, which sets the scale of the x-axis. Red line and shaded fill indicate the mean and s.e.m. of nonlinearity thresholds (see text for details). (c) Visualization of the principal axes of variation in subunit nonlinearities by adding or subtracting principal components from the mean nonlinearity. (top) The principal axis of variation in subunit nonlinearities results in a gain change, while (bottom) the second principal axis corresponds to a threshold shift. These two dimensions captured 63% of the nonlinearity variability across cells.

<https://doi.org/10.1371/journal.pcbi.1006291.g005>

factor: i.e.  $25 \rightarrow \sqrt{25} = 5$ . This estimate is largely consistent with the typical number of subunits in Fig 5A, required to optimize model predictive performance. This estimate also suggests that the large majority of bipolar-to-ganglion cell synapses are rectifying (strongly nonlinear), as linear connections are not uniquely identifiable in an LN-LN cascade. Indeed, in the salamander retina, strong rectification appears to be the norm [45, 50].

**LN-LN models have subunit nonlinearities with high thresholds.** The nonlinearities for all of the measured subunits are overlaid in Fig 5B. Each LN-LN subunit nonlinearity takes as input the projection of the stimulus onto the corresponding subunit spatiotemporal filter. Since the stimulus components are white noise with unit standard deviation, and the spatiotemporal filter is constrained to have unit norm, the projection of the stimulus onto the filter has a standard Normal distribution, thus we can compare nonlinearities on a common axis. Despite the fact that the model could separately learn an arbitrary function over the input for each subunit nonlinearity, we find that the nonlinearities are fairly consistent across the different subunits of many cells. Subunit nonlinearities look roughly like thresholding functions, relatively flat for most inputs but then increasing sharply after a threshold. We quantified the threshold as input for which the nonlinearity reaches 40% of the maximum output, across  $n = 92$  model-identified subunits the mean threshold was  $3.09 \pm 0.14$  (s.e.m.) standard deviations. We additionally computed LN thresholds for ganglion cells and found that they were similarly consistent across the population ( $3.22 \pm 0.11$  standard deviations). We decomposed the set of nonlinearities using principal components analysis and show the two primary axes of variation in Fig 5C. The primary axis of variation results in a gain change, while the secondary axis induces a threshold shift. Due to the high thresholds of these nonlinearities, subunits only impact ganglion cell firing probability for large input values. The slight rise on the left side of the nonlinearities is likely due to weak ON- inputs to the ganglion cell, and a stimulus ensemble that drives the ON- pathways more strongly [39, 55] may be necessary to uniquely identify them. Our ganglion cell population consisted of 18 Fast-Off, 4 Slow-Off, and 1 On cell (classification shown in S2 Fig). We did not find significant differences across cell types in terms of the number of identified subunits or the subunit thresholds.

## Computational properties of learned LN-LN models


We now turn from a quantitative analysis of the physiological properties of the retina, described above, to their implications in terms of the computational function of the retina in processing visual stimuli. In particular, in the next two sub-sections we predict that the dominant contribution to stimulus decorrelation in efficient coding theory occurs at the bipolar cell synaptic threshold, and that the composite function computed by a retinal ganglion cell corresponds to a logical OR of its bipolar cell inputs.

**Stimulus decorrelation at different stages in hierarchical retinal processing.** Natural stimuli have highly redundant structure. Efficient coding theories [56, 57] state that sensory systems ought to remove these redundancies in order to efficiently encode natural stimuli. The simplest such redundancy is that nearby points in space and time contain similar, or correlated, luminance levels [58]. The transmission of such correlated structure would thus be highly inefficient. Efficient coding has been used to explain why responses are much less correlated than natural scenes, although the mechanistic underpinnings of decorrelation in the retina remain unclear.

Early work [59] suggested a simple mechanism: the linear center-surround receptive field of ganglion cells (and more recently, of bipolar cells [60]) could contribute to redundancy reduction simply by transmitting only differences in stimulus intensity across nearby positions in space. However, it was recently shown [61] using LN models that most of the decorrelation

of naturalistic stimuli in the retina could be attributed to ganglion cell nonlinearities, as opposed to linear filtering. Given that we fit an entire layer of subunits pre-synaptic to each ganglion cell layer, we can analyze the spatial representation of naturalistic images at different stages of hierarchical retinal processing, thereby localizing the computation of decorrelation to a particular stage in the model.

To do so, we generated the response of the entire population of model subunits to a spatial stimulus similar to previous work [61], namely spatially pink noise, low pass filtered in time. We computed the correlation of stimulus intensities as a function of spatial distance, as well as the correlation between pairs of model units as a function of spatial distance. We examined pairs of units across different stages of the LN-LN model: after linear filtering by the subunits, after the subunit nonlinearity, and finally at the ganglion cell firing rates (the final stage). Fig 6 shows that the correlation at these stages drops off with distance between either the subunits or ganglion cells (with distance measured between receptive field centers). Pitkow & Meister [61] found that the nonlinearity in an LN model was primarily responsible for most of the decorrelation (overall suppression of the correlation curve towards zero). However, without the ability to map the LN nonlinearity to biophysical components in the retina, it is hard to map this result onto a particular biophysical mechanism. Instead, our model predicts that this decorrelation is primarily due to the *subunit* nonlinearities, as opposed to ganglion cell spiking nonlinearities. In fact, the correlation between the ganglion cell model firing rates slightly increases after pooling across subunits. The most decorrelated representation occurs just after thresholding at the subunit layer. In this manner, our modeling framework suggests a more precisely localized mechanistic origin for a central tenet of efficient coding theory. Namely, our results predict that the removal of visual redundancies, through stimulus decorrelation



**Fig 6. Decorrelation in LN-LN subunit models.** A naturalistic (pink noise) stimulus was shown to a population of nonlinear subunits. The correlation in the population after filtering at the subunit layer (blue), after the subunit nonlinearity (green), and after pooling and thresholding at the ganglion cell layer (red), in addition to the stimulus correlations (gray) are shown. Left: the correlation as a function of distance on the retina for Off-Off cell pairs. Right: correlation for Off-On cell pairs. For each plot, distances were binned every 70 $\mu$ m, and error bars are the s.e.m. within each bin.

<https://doi.org/10.1371/journal.pcbi.1006291.g006>


across space, originates primarily from high-threshold nonlinearities associated with bipolar cell synapses.

**The nature of nonlinear spatial integration across retinal subunits.** Retinal ganglion cells emit responses in sparse, temporally precise patterns [62], presumably to keep firing rates low thereby preserving energy [61, 63]. LN models can emulate sparse, precise firing in only one way: by using nonlinearities with high thresholds relative to the distribution of stimuli projected onto their linear filter. This way, only a small fraction of stimuli will cause the model to generate a response. Indeed, nonlinearities in LN models fit to ganglion cells have high thresholds [9]. LN-LN models, in contrast, can generate sparse responses using two qualitatively distinct operating regimes: either the subunit thresholds (first nonlinearity) could be high and the ganglion cell or spiking threshold (second nonlinearity) could be low, or the subunit thresholds could be low and spiking thresholds could be high. Both of these scenarios give rise to sparse firing at the ganglion cell output. However, they correspond to categorically distinct functional computations.

These various scenarios are diagrammed in Fig 7A–7C. Each panel shows the response of a model in a two-dimensional space defined by the projection of the stimulus onto two subunit filters pre-synaptic to the ganglion cell (the two-dimensional space is easier for visualization, but the same picture holds for multiple subunits). We show the response as contours where the firing probability is constant (iso-response contours). Here, the subunit nonlinearities play a key role in shaping the geometry of the response contours, and therefore shape the computation performed by the cell. Note that the ganglion cell nonlinearity would act to rescale the output, but cannot change the shape of the contours. Therefore, it is fundamentally the subunit nonlinearities alone that determine the geometry of the response contours.

Low-threshold subunit nonlinearities give rise to concave contours (Fig 7B), whereas high-threshold subunits give rise to convex contours (Fig 7C). Because final output rate is determined by the subunit and final thresholds, both of these descriptions could yield sparse firing output with the same overall rate (by adjusting the final threshold), but correspond to different computations. Low-threshold subunits can be simultaneously active across many stimuli, and thus yield spiking when subunits are simultaneously active (an AND-like combination of inputs). On the other hand, high-threshold subunits are rarely simultaneously active and thus usually only one subunit is active during a ganglion cell firing event, giving rise to an OR-like combination of inputs. By comparison, a cell that linearly integrates its inputs would have linear contours (Fig 7A).

In our models fit to retinal ganglion cells, we find all cells are much more consistent with the high threshold OR-like model. Subunit nonlinearities tend to have high thresholds, and therefore result in convex contours (shown for different pairs of subunits for two example cells in Fig 7D and 7E). For each example ganglion cell, we show the corresponding model contours along with the 2 standard deviation contour of the stimulus distribution (gray oval) and the empirical firing histogram (red checkers) in the 2D space defined by the projection of the stimulus onto a given pair of subunit filters identified by the LN-LN cascade model. Note that while the stimulus is uncorrelated (i.e. white, or circular), non-orthogonality of subunit filters themselves yield correlations in the subunit activations obtained by applying each subunit filter to the stimulus. Hence the stimulus distribution in the space of subunit activations (grey shaded ovals) is not circular. In all recorded cells, we find that the composite computation implemented by retinal ganglion cell circuitry corresponds to an OR function associated with high subunit thresholds (as schematized in Fig 7C). Moreover, both the AND computation and the linear model are qualitatively ruled out by the shape of the model response contours as well as the empirical firing histogram over subunit activations, which closely tracks the model response contours (i.e. the boundaries of the red histograms are well captured by the model contours).



**Fig 7. Visualization of subunit contours.** Contours of equal firing probability are shown in a 2D space defined by the projection of the visual stimulus along each of two subunits. (a) Example contour plots for a model with low threshold subunit nonlinearities (inset) has concave contours. (b) A model with high threshold subunit nonlinearities has convex contours. (c & d) Contours from a model for two example ganglion cells, for three different pairs of subunits (left to right). In each panel, a histogram of the recorded firing rate is shown (red squares) as well as the stimulus distribution (gray oval).


<https://doi.org/10.1371/journal.pcbi.1006291.g007>

These results are consistent with previous studies of nonlinear spatial integration in the retina. For example, Bollinger et al. [45] discovered convex iso-response contours for a very simple two dimensional spatial stimulus, and Kaardal et al. [44] performed an explicit hypothesis test between an AND-like and OR-like nonlinear integration over a low dimensional subspace obtained via the un-regularized STC eigenvectors, finding that OR outperformed AND. However, the techniques of [45] can only explore a low-dimensional stimulus space, whereas our methods enable the discovery of iso-response contours for high-dimensional stimuli.

Moreover, in contrast to the hypothesis testing approach taken in [44], our general methods to learn LN-LN models reveal that an OR-model of nonlinear integration is a good model on an absolute scale of performance amongst all models in the LN-LN family, rather than simply being better than an AND-model.

**A multi-dimensional view of cascaded retinal computation.** A simple, qualitative schematic of the distinct computational regime in which retinal ganglion cells operate in response to white noise stimuli can be obtained by considering the geometry of the spike triggered ensemble in  $N$  dimensional space. In particular, the distribution of stimuli concentrates on a constant radius sphere in  $N$  dimensional stimulus space. More precisely, any high dimensional random stimulus realization  $\mathbf{x}$  has approximately the same vector length, because the fluctuations in length across realizations, relative to the mean length, is  $O(1/\sqrt{N})$ . Thus we can think of all likely white-noise stimuli as occurring on the  $N - 1$  dimensional surface of a sphere in  $N$  dimensional space. Each subunit filter can be thought of as a vector pointing in a particular direction in  $N$  dimensional stimulus space. The corresponding input to the subunit nonlinearity for any stimulus is the inner-product of the stimulus with the subunit filter, when both are viewed as  $N$  dimensional vectors. The high threshold of the subunit nonlinearity means that the subunit only responds to a small subset of stimuli on the sphere, corresponding to a small cap centered around the subunit filter. For a single subunit model (i.e. an LN model), the set of stimuli that elicit a spike then corresponds simply to this one cap (Fig 8A). In contrast, the OR like computation implemented by an LN-LN model with high subunit thresholds responds to stimuli in a region consisting of a union of small caps, one for each subunit (Fig 8B).

To verify this conceptual picture, in Fig 8C we visualize a two-dimensional projection of the spike-triggered stimulus ensemble for three example ganglion cells, using principal components analysis of the spike-triggered subunit activations. That is, we project the spike-triggered ensemble onto the subunit filters identified in the LN-LN model, and subsequently project those subunit activations onto the two dimensions that capture the most variance in subunit activations. Note that this is different from just taking the top two principal components of the STC matrix, as the top STC component is typically the *average* of the subunit filters [10], which does not differentiate the subunit activations. This procedure identifies a subspace that captures the radial spread of subunit filters in high-dimensional stimulus space. We find that



**Fig 8. Stimulus selectivity in LN and LN-LN models.** Each panel shows the raw stimulus distribution (gray contours) projected onto the top two principal components of the spike-triggered subunit activations (with subunits identified by the LN-LN model). The LN model (a) fires in response to stimuli in a single region, or cap, of stimulus space (indicated by the arrow and dashed threshold), whereas the LN-LN model (b) fires in response to a union of caps, each defined by an individual subunit. (c) Spike-triggered subunit activations for three representative cells are shown as colored histograms (colors indicate which model-identified subunit was maximally active during the spike), with the corresponding subunit filter directions shown as colored arrows (see text for details). Color intensity of the histogram indicates the probability density of the spike-triggered ensemble (STE), thus drops in intensity between changes in color indicate a multimodal STE, with high density modes centered near subunit filter directions.

<https://doi.org/10.1371/journal.pcbi.1006291.g008>

the spike-triggered ensemble projected onto this subspace (Fig 8C) curves around the radial shell defined by the stimulus distribution, and matches the conceptual picture shown in Fig 8B. For ease of visualization, we colored elements in the spike-triggered ensemble by which LN-LN model subunit was maximally active during that spike, and we normalize the spike-triggered histogram by the raw stimulus distribution (gray ovals). This picture provides a simple, compelling view for why LN models are insufficient to capture the retinal response to white noise, and further visualizes the aspect of retinal computation LN-LN models capture that the LN model does not: ganglion cells encode the union of different types of stimuli, with each stimulus type having large overlap with precisely one subunit filter.


### Regularized spike-triggered analysis

Given a mathematical model of a multilayered neural circuit, we can connect the pathways in such models back to descriptive statistics, namely, spike-triggered statistics such as the spike-triggered average (STA) and covariance (STC). That is, we can show that the STA and the STC eigenvectors of a general LN-LN model are linear combinations its pathway filters (see [Methods](#) and [64]). Therefore, we expect certain types of structure in said pathways to persist after the linear combination, assuming the number of pathways is small relative to the stimulus dimension. This immediately suggests that the same proximal algorithms and regularization terms we used to fit LN-LN models can be used to regularize the STA and the STC eigenvectors directly (for situations where one is interested in the descriptive statistics, but not the full encoding model).

To illustrate the benefits of the regularization terms used to fit the LN-LN models, we apply these penalties to perform regularized spike-triggered analysis. We formulate optimization problems for regularizing the spike-triggered average and covariance which only require access to the un-regularized estimates (see [Methods](#)). This is useful for the situation where working with the full spike-triggered ensemble or raw dataset is prohibitive due to computational time or memory constraints.

[Fig 9A](#) compares a regularized spike-triggered average with the raw, un-regularized STA for an example recorded ganglion cell in response to a 1-D spatiotemporal white noise stimulus. For long recordings, the regularized STA closely matches the raw STA, while for short recordings the regularized STA has less high frequency noise and retains much of the structure observed if the STA had been estimated using more data. [Fig 9B](#) shows the held-out performance of the regularized STA for an example cell across different regularization weights, scanned over a broad range, demonstrating that performance is largely insensitive to the strengths of the weights of the  $\ell_1$  and nuclear norm penalty functions. Thus regularization weights need not be fine tuned to achieve superior performance. We further quantified the performance of the regularized STA by using it as the linear filter of an LN model, and found that with regularization, about 5 minutes of recording was sufficient to achieve the performance (on held-out data) obtained through 40 minutes of recording without regularization ([Fig 9C](#)).

[Fig 10](#) demonstrates the improvement in our ability to estimate the relevant subspace spanned by significant STC eigenvectors, both in terms of the qualitative improvement in eigenvectors for an example cell ([Fig 10A](#)) and quantified across the population ([Fig 10B](#)). In [Fig 10A](#), we show the top regularized STC eigenvectors for different values of the nuclear norm ( $\gamma_n$ ) and  $\ell_1$ -norm ( $\gamma_1$ ) regularization penalties (Eq 11 in [Methods](#) and [Table 1](#)). We score the performance of the STC subspace in [Fig 10B](#) in terms of how well stimuli, after projection onto the subspace, can be used to predict spikes, by computing the subspace overlap (defined in [Methods](#)) between the raw or regularized STC subspace and the best fit LN-LN subspace.



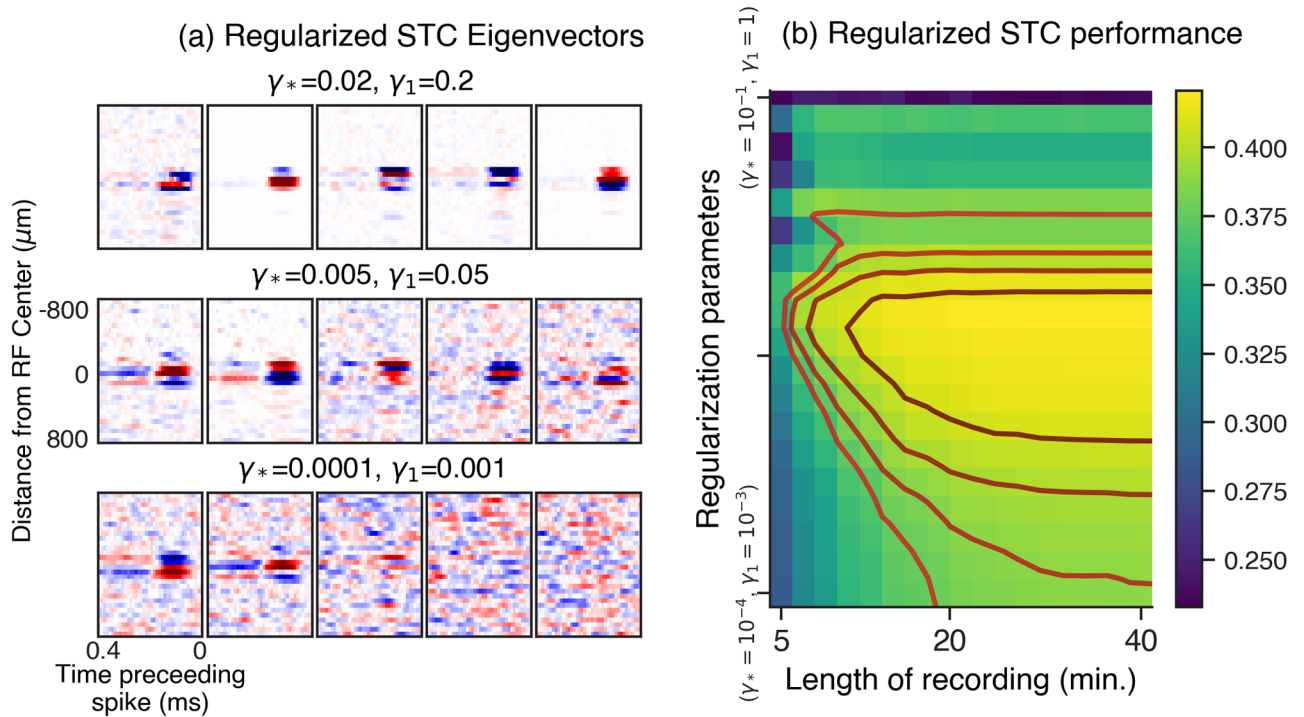
**Fig 9. Regularization for estimating receptive fields (via a regularized spike-triggered-average).** (a) Top row: the raw spike-triggered average computed using different amounts of data (from left to right, 30s to 40min), bottom row: the regularized spike-triggered average computed using the same amount of data as the corresponding column. (b) Performance (held-out log-likelihood) as a function of two regularization weights, the nuclear norm (x-axis, encourages low-rank structure) and the  $\ell_1$ -norm (y-axis, encourages sparsity), for an example cell. (c) Correlation coefficient (on held-out data) between the firing rate of a retinal ganglion cell and LN model whose filter is fixed to be a regularized or raw (un-regularized) STA, as a function of the amount of training data for estimating the STA (length of recording).

<https://doi.org/10.1371/journal.pcbi.1006291.g009>

This quantity ranges between zero for orthogonal subspaces and one for overlapping subspaces.

Since the LN-LN subspace is the best subspace found by the LN-LN model for predicting spiking, a large subspace overlap between the regularized STC and LN-LN subspaces indicates the ability of regularized STC to find stimulus subspaces predictive of neural firing without actually fitting a model of the neuron. With appropriate regularization, one can recover the best predictive subspace using about 10 minutes of data; without regularization, one requires 40 minutes of data to recover a subspace with comparative predictive accuracy. Note that even for the full length of this experiment (40 minutes), regularization still improves our regularized STC estimate. Thus, we find the subspace spanned by regularized STC eigenvectors becomes very similar to the subspace obtained from the filters in an LN-LN model, as predicted by our theoretical analysis.





**Fig 10. Regularized spike-triggered covariance.** (a) Example panels of the output of our regularized spike-triggered covariance algorithm. Each panel contains the five most significant regularized eigenvectors of the STC matrix, reshaped as spatiotemporal filters. The bottom panel shows the result with no regularization added, and the upper panels show the result with increasing weights on the regularization penalties. Here  $\gamma_1$  is the regularization weight applied to an  $\ell_1$  penalty encouraging sparsity, and  $\gamma^*$  is a regularization weight applied to a nuclear norm penalty, encouraging approximate spatiotemporal separability of the eigenvectors, when reshaped as spatiotemporal filters. (b) Summary across a population of cells. The heatmap shows the held-out performance of regularized STC (measured as the subspace overlap with the best fit LN-LN subspace, see text for details). The y-axis in (b) represents a line spanning 3 orders of magnitude in two-dimensional regularization parameter space  $(\gamma^*, \gamma_1)$ , ranging from the point  $(\gamma^* = 10^{-4}, \gamma_1 = 10^{-3})$  to  $(\gamma^* = 10^{-1}, \gamma_1 = 1)$ .

<https://doi.org/10.1371/journal.pcbi.1006291.g010>

## Discussion

In summary, we combine proximal algorithms with non-smooth regularization terms to model stimulus driven neural processing in circuits with multiple parallel, hierarchical nonlinear pathways using limited experimental data. We found that models employing two stages of

**Table 1. Common regularization penalties and their proximal operators (in closed form).**

Penalty function, $\phi(x)$	Proximal operator, $\mathcal{P}_\phi(v, \rho)$	Computational complexity
$\ell_2$ -norm, $\gamma\ x\ _2$	$\frac{v}{1+\gamma/\rho}$	$\mathcal{O}(n)$
$\ell_1$ -norm, $\gamma\ x\ _1$	$\begin{cases} v_i - \gamma/\rho & v_i \geq \gamma/\rho \\ v_i + \gamma/\rho & v_i \leq -\gamma/\rho \\ 0 & \text{otherwise} \end{cases}$	$\mathcal{O}(n)$
Nuclear norm, $\ X\ _*$	$\mathbf{X} = \mathbf{USV}^T, \mathcal{P}_\phi(v, \rho) = \mathbf{US}'\mathbf{V}^T,$ $\mathbf{S}' = \begin{cases} \sigma_i - \gamma/\rho & \sigma_i \geq \gamma/\rho \\ 0 & \text{otherwise} \end{cases}$	$\mathcal{O}(\min(mn^2, nm^2))$
Non-negativity, $\mathcal{I}(x > 0)$	$\begin{cases} v_i & v_i > 0 \\ 0 & v_i \leq 0 \end{cases}$	$\mathcal{O}(n)$

<https://doi.org/10.1371/journal.pcbi.1006291.t001>

linear and nonlinear computation, namely LN-LN models, demonstrated a robust improvement over the classical standard of LN models at predicting responses to white noise across a population of ganglion cells. Beyond performance considerations alone, the gross architecture of the LN-LN model maps directly onto the hierarchical, cascaded, anatomy of the retina, thereby enabling the possibility that we can generate quantitative hypotheses about neural signal propagation and computation in the unobserved interior of the retina simply by examining the structure of our model's interior. Since learning our model only requires measurements of the inputs and outputs to the retinal circuit, this approach is tantamount to the computational reconstruction of unobserved hidden layers of a neural circuit. The advantage of applying this method in the retina is that we can experimentally validate aspects of this computational reconstruction procedure.

Indeed, using intracellular recordings of bipolar cells, we found that our learned subunits matched properties of bipolar cells, both in terms of their receptive field center-surround structure, and in terms of the approximate number of bipolar cells connected to a ganglion-cell. However care must be taken not to directly identify the learned subunits in our model with bipolar cells in the retina. Instead, they should be thought of as functional subunits that reflect the combined contribution of not only bipolar cells, but also horizontal cells and amacrine cells that sculpt the composite response of retinal ganglion cells to stimuli. Nevertheless, the correspondence between subunits and bipolar cell RFs (which are also shaped by horizontal cells), suggests that even for Gaussian white noise stimuli, it is important to learn functional subunits that loosely correspond to the composite effect that bipolar cells and associated circuitry have on ganglion-cell responses.

The interior of our models also reveal several functional principles underlying retinal processing. First, all subunits across all cells had strikingly consistent nonlinearities corresponding to monotonically increasing threshold-like functions with very high thresholds. This inferred biophysical property yields several important consequences for neural signal processing in the inner retina. First, it predicts that subunit activation patterns are sparse across the ensemble of stimuli, with typically only one subunit actively contributing to any given ganglion cell spike. Second, it predicts that the dominant source of stimulus decorrelation, a central tenet of efficient coding theory, has its mechanistic origin at the first strongly nonlinear processing stage of the retina, namely in the synapse from bipolar cells to ganglion cells. Third, it implies that the composite function computed by individual retinal ganglion cells corresponds to a Boolean OR function of bipolar cell feature detectors.

Taken together, the proximal algorithms framework provides a unified way to both estimate hierarchical nonlinear models of sensory processing and compute spike-triggered statistics using limited data. When applied to the retina, these techniques recover aspects of the interior of the retina without requiring direct measurements of retinal interneurons. Moreover, by identifying candidate mechanisms for cascaded nonlinear computation in retinal circuitry, our results provide a higher resolution view of retinal processing compared to classic LN models, thereby setting the stage for the next generation of efficient coding theories that may provide a normative explanation for such processing. For example, considerations of efficient coding have been employed to explain aspects of the linear filter [65] and nonlinearity [61] of retinal ganglion cells when viewed through the coarse lens of an LN model. An important direction for future research would be the extension of these basic theories to more sophisticated ones that can explain the higher resolution view of retinal processing uncovered by our learned LN-LN models. Principles that underlie such theories of LN-LN processing might include subthreshold noise rejection [66, 67], sensitivity to higher order statistical structure in natural scenes, and energy efficiency [63]. Indeed the ability to extract these models

from data in both a statistically and computationally efficient manner constitutes an important step in the genesis and validation of such a theory.

Another phenomenon robustly observed in the retina is adaptation to the luminance and contrast of the visual scene. Adaptation is thought to be a critical component of the retinal response to natural scenes [30], and a promising direction for extensions of our work would be to include luminance and contrast adaptation in subunit models. Luminance adaptation (adapting to the mean light intensity) is mediated by photoreceptor cells, and could be modeled by prepending a simple photoreceptor model (e.g. [68]) to an LN-LN model. There are two major sites of contrast adaptation, at the bipolar-to-ganglion cell synapse [69, 70] and at the spiking mechanism of ganglion cells [69, 71]. Extending the simple thresholding nonlinearities in our model with a dynamical model of adaptation (e. g. [19]) is a first step towards understanding the interaction between nonlinear subunits and adaptation.

While our work utilized white noise stimuli, the methods do not require any particular form of stimulus and will thus generalize to other stimulus distributions. In particular, stimuli that differentially activate subunits will be the most effective at differentiating LN and LN-LN models. Stimuli with coarse spatial resolution will not differentially activate subunits within the receptive field, thus are a poor choice for studying nonlinear spatial integration. However, fine textures as present in natural stimuli, are very likely to activate these nonlinear mechanisms in the retina, and thus are a critical component for understanding vision in the context of ethologically relevant stimuli.

The computational motifs identified by LN-LN models are likely to generalize across different species because they rely on a few key properties. For example, our predictions about the primary source of decorrelation in the retina rely on three features of the underlying circuitry identified by LN-LN models: (a) bipolar cell receptive fields are smaller than those of ganglion cells, (b) bipolar cell receptive field centers are largely non-overlapping, and (c) bipolar cell synapses have high thresholds. In addition, the logical OR combination of features relies on high thresholds and bipolar receptive fields that are (largely) non-overlapping. These properties (high threshold subunits with smaller, non-overlapping receptive fields) are common across multiple species.

Beyond the retina, multiple stages of cascaded nonlinear computation constitutes a ubiquitous motif in the structure and function of neural circuits. The tools we have applied here to elucidate hierarchical nonlinear processing in the retina are similarly applicable across neural systems more generally. Thus we hope our work provides mathematical and computational tools for efficiently extracting and analyzing both informative descriptive statistics and hierarchical nonlinear models across many different sensory modalities, brain regions, and stimulus ensembles, thereby furthering our understanding of general principles underlying nonlinear neural computation.

## Methods

### Ethics statement

This study was performed in strict accordance with the recommendations in the Guide for the Care and Use of Laboratory Animals of the National Institutes of Health, and the Stanford institutional animal care and use committee (IACUC) protocol (11619).

### Experiments

Experimental data was collected from the tiger salamander retina using a multi-electrode array (Multi-Channel Systems), as described elsewhere [9]. Isolated ganglion cells were identified using custom spike sorting software. The stimulus used was a 100 Hz white noise bars stimulus,

where the luminance of each bar was drawn independently from a Gaussian distribution. Spatially, the stimulus spanned approximately 2.8 mm on the retina (50 bars at 55.5  $\mu\text{m}$  / bar). Intracellular recordings were performed as described elsewhere [72]. Off bipolar cells were identified by their flash response, receptive field size, and level in the retina. Data were analyzed using the `pyret` package [73]. The experimental data, as well as general purpose code to fit hierarchical LN-LN models and perform regularized STA and STC analysis, are available at <https://github.com/baccuslab/inferring-hidden-structure-retinal-circuits>.

### LN-LN models

In this section, we specify the mathematical formulation of our LN-LN models. The model takes a spatiotemporal stimulus, represented as a vector  $\mathbf{x}$ , and generates a predicted firing rate,  $r(\mathbf{x})$ . First, the stimulus is projected onto a number of subunit filters. The number of subunits is a hyper-parameter of the model, chosen through cross validation (we repeatedly fit models with increasing numbers of subunits until held-out performance on a validation set decreases). If we have  $k$  subunits, then the stimulus is projected onto each of the  $k$  filters:  $\mathbf{w}_i$  for  $i = 1, \dots, k$ . These projections are then passed through separate subunit nonlinearities. The nonlinearities are parameterized using a set of Gaussian basis functions (or bumps) that tile the relevant input space [12, 20]. This parameterization enforces smoothness of the nonlinearity. We typically use  $p = 30$  evenly spaced Gaussian bumps that tile the range spanned by the projection of the stimulus onto the linear filter (results were not sensitive to the number of bumps over a range of 10–30 bumps). For example, a nonlinearity  $h(u)$  is parameterized as

$$\begin{aligned} h(u) &= \sum_{j=1}^p a_j \phi_j(u) \\ &= \sum_{j=1}^p a_j \phi(u - \Delta_j), \end{aligned}$$

where  $\phi$  is the basis function, e.g.  $\phi(x) = \exp(-x^2)$ ,  $\Delta_j$  indicates the spacing between the basis functions,  $p$  is the number of bases used, and  $a_j$  is a weight on that particular basis function. Since the basis functions and spacings are fixed beforehand, the only free parameters are the  $a_j$ 's. For subunit  $i$ , the corresponding nonlinearity has a set of weights  $a_{ij}$  for  $j = 1, \dots, p$ .

The output of the  $k$  subunits is then summed and passed through a final nonlinearity. This final nonlinearity is parameterized as a soft rectifying function  $r(x) = g \log(1 + e^{x-\theta})$ , with two parameters:  $g$  is an overall gain and  $\theta$  is the threshold. The full LN-LN model is then given by:

$$r(\mathbf{x}) = g \log \left( 1 + \exp \left[ \left( \sum_{i=1}^k \sum_{j=1}^p a_{ij} \phi(\mathbf{w}_i^T \mathbf{x} - \Delta_j) \right) - \theta \right] \right),$$

where the parameters to optimize are the subunit filters  $\mathbf{w}_i$  for  $i = 1, \dots, k$ , subunit nonlinearity weights  $a_{ij}$  for  $j = 1, \dots, p$ , and final nonlinearity parameters  $\theta$  and  $g$ .

We optimize the parameters using a maximum likelihood objective assuming a Poisson noise model for spiking. Rather than optimize all of the parameters simultaneously, we alternate between optimizing blocks of parameters (joint optimization using gradient descent was prone to getting stuck at solutions that were less accurate). That is, we alternate between optimizing three blocks of parameters: the subunit filters  $\mathbf{w}_i$ , the subunit nonlinearities  $a_{ij}$ , and the final nonlinearity parameterized by  $\theta$  and  $g$ . We optimize each block of parameters by minimizing the negative log-likelihood of the data plus any regularization terms using proximal algorithms. The subunit filters are the only parameters with regularization

penalties (the nuclear norm applied to the filter reshaped as a spatiotemporal matrix and the  $\ell_1$  norm), to encourage space-time separability and sparseness of the filters. The proximal operator for each of these regularization penalties is given in Table 1, and the proximal operator for the log-likelihood term (which does not have a closed-form solution) is solved using gradient descent. In addition, after optimizing the block of parameters corresponding to the subunit filters, we rescale them to have unit norm before continuing the alternating minimization scheme. This ensures that the distribution of input to the nonlinearities spans the same range, and gets rid of an ambiguity between the scale of the subunit filters and the scale of the domain of the subunit nonlinearity. We find that the parameters converge after several rounds of alternating minimization, and are robust with respect to random initialization of the parameters.

### Proximal operators and algorithms

The framework of *proximal algorithms* allows us to efficiently optimize functions with non-smooth terms. The name *proximal* comes from the fact that these algorithms utilize the *proximal operator* (defined below) as subroutines or steps in the optimization algorithm. For brevity, we skip the derivation of these algorithms, instead referring the reader to the more thorough treatment by Parikh and Boyd [74] or Polson et al. [75]. The proximal operator for a function  $\phi$  given a starting point  $v$  is defined as:

$$\mathcal{P}_\phi(v) = \operatorname{argmin}_x \left[ \phi(x) + \frac{\rho}{2} \|x - v\|_2^2 \right]. \quad (1)$$

The proximal operator is a mapping from a starting point  $v$  to a new point  $x$  that tries to minimize the function  $\phi(x)$  (first term above) but stays close to the starting point  $v$  (second term), where the parameter  $\rho$  trades off between these two objectives. The proximal operator is a building block that we will use to create more complicated algorithms. We will take advantage of the fact that for many functions  $\phi$  of interest to us, we can analytically compute their proximal operators, thus making these operators a computationally cheap building block.

We used these building blocks to solve optimization problems involving the sum of a number of simpler terms:

$$\min_x \sum_{i=1}^k \phi_i(x) \quad (2)$$

where in our application the  $\phi_i$ 's represent either a model fitting objective (e.g. a log-likelihood) or different regularization penalties on the parameters,  $x$ . For example, for learning the parameters of a linear filter in an LN model, the objective consists of a log-likelihood  $f(x)$  along with regularization penalties that impose prior beliefs about the filter,  $x$ . We focus on two main penalties. Sparsity, which encodes the belief that many filter coefficients are zero, is penalized by the  $\ell_1$ -norm ( $\phi_1(x) = \|x\|_1$ ). Additionally, spatiotemporal filters are often approximately space-time separable (they are well modeled as the outer product of a few spatial and temporal factors). We encoded this penalty by the nuclear norm,  $\ell_*$ , which encourages the parameters  $x$ , when reshaped to form a spatiotemporal matrix, to be a low-rank matrix (the nuclear norm  $\ell_*$  of a matrix is simply the sum of its singular values). Another natural penalty would be one that encourages the parameters to be smooth in space and/or time, which could be accomplished by applying an  $\ell_1$  or  $\ell_2$  penalty to the spatial or temporal differences in parameters. As shown below, these types of penalties are easy to incorporate into the proximal algorithm framework. Other commonly used regularization penalties, and their corresponding proximal operators, are listed in Table 1.

The proximal consensus algorithm is an iterative algorithm for solving (2) that takes a series of proximal operator steps. It first creates a copy of the variable  $x$  for each term  $\phi_i$  in the objective. The algorithm proceeds by alternating between taking proximal steps for each function  $\phi_i$  using that variable copy  $x_i$ , and then enforcing all of the different variable copies to agree (reach consensus) by averaging them. The algorithm is:

$$\begin{aligned} x_i^{k+1} &= \mathcal{P}_{\phi_i}(\bar{x}^k - u_i^k) \\ \bar{x}^{k+1} &= \frac{1}{k} \sum_{i=1}^k x_i \\ u_i^{k+1} &= u_i^k + x_i^{k+1} - \bar{x}^{k+1}, \end{aligned}$$

where  $i$  indexes each of the terms in the objective function,  $x_i$  is a copy of the variable,  $\bar{x}$  is the average of the variable copies, and  $u_i$  is a dual variable that can be thought of as keeping a running average of the error between each variable copy and the average. Intuitively, we can think of each variable copy  $x_i$  as trying to minimize a single term  $\phi_i$  in the objective, and the average, or consensus  $\bar{x}$  forces the different copies to agree. After convergence, each copy  $x_i$  will be close to the mean value  $\bar{x}$ , which is the set of parameters that minimizes the original composite objective.

This algorithm has a number of desirable properties. First, the updates for each term  $x_i$  can be carried out in parallel, therefore allowing for speedups when run on a cluster or multi-core computer. Second, it converges even when terms in the objective are non-differentiable. Due to the repeated application of the proximal operator, this algorithm works best when the terms  $\phi_i$  have proximal operators that are easy to compute.

This is exactly the case for the regularization terms described above: for the  $\ell_1$  norm, the proximal operator corresponds to soft thresholding of the parameters. For the nuclear norm, the proximal operator corresponds to soft thresholding of the singular values of parameters reshaped as a matrix. Occasionally, the proximal operator may not have a closed form solution. In this case, the proximal step can be carried out through gradient based optimization of (1) directly. This is the case for some log-likelihoods, such as the log-likelihood of a particular firing rate under Poisson spiking. In this case, gradient step based optimization of (1) often dominates the computational cost of the algorithm. As many methods for fitting neural models involve gradient step updates on the log-likelihood, such methods can then be augmented with additional regularization terms with no appreciable effect on runtime, by using proximal consensus algorithms for optimization. Our code for solving formulating and solving optimization problems using proximal algorithms is provided online at <https://github.com/ganguli-lab/proxalgs>.

## Relationship between descriptive statistics and encoding models

Here, we derive the relationship between the pathways of any differentiable encoding model and spike-triggered statistics under Gaussian noise stimulation. We represent a visual stimulus as an  $N$  dimensional vector  $\mathbf{x}$ . We view a functional neural model as an arbitrary nonlinear function  $r = f(\mathbf{x})$ , over  $N$  dimensional stimulus space, where  $r$  determines the probability that the neuron fires in a small time window following a stimulus  $\mathbf{x}$ :  $r(\mathbf{x}) = p(\text{spike} | \mathbf{x})$ . The derivation will show how the STA is related to the *gradient* of the model  $\nabla r(\mathbf{x})$ , and the STC is related to the *Hessian*,  $\nabla^2 r(\mathbf{x})$ .

The STA and STC are the mean and covariance, respectively, of the spike-triggered stimulus ensemble, which reflects the collection of stimuli preceding each spike [3]. This distribution

over stimuli, conditioned on a spike occurring, can be expressed via Bayes rule,

$$p(\mathbf{x} \mid \text{spike}) = \frac{p(\text{spike} \mid \mathbf{x})p(\mathbf{x})}{p(\text{spike})}, \tag{3}$$

where  $p(\mathbf{x})$  is the prior distribution over stimuli and  $p(\text{spike})$  is the average firing probability over all stimuli. Here, we assume a white noise stimulus distribution, in which each component of  $\mathbf{x}$  is chosen independently from a Gaussian distribution with zero mean and unit variance. The STA and STC are given by

$$\mathbf{x}_{\text{STA}} = \mathbb{E}_{p(\mathbf{x}|\text{spike})}[\mathbf{x}] \tag{4}$$

$$\mathbf{C}_{\text{STC}} = \mathbb{E}_{p(\mathbf{x}|\text{spike})}[\mathbf{x}\mathbf{x}^T] - (\mathbf{x}_{\text{STA}})(\mathbf{x}_{\text{STA}})^T \tag{5}$$

Focusing first on the STA:

$$\begin{aligned} \mathbf{x}_{\text{STA}} &= \int \mathbf{x} p(\mathbf{x} \mid \text{spike}) d\mathbf{x} \\ &= \frac{1}{\mu} \int \mathbf{x} r(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \\ &= \frac{1}{\mu} \mathbb{E}_{p(\mathbf{x})}[\mathbf{x} r(\mathbf{x})] \\ &= \frac{1}{\mu} \mathbb{E}_{p(\mathbf{x})}[\nabla r(\mathbf{x})], \end{aligned} \tag{6}$$

where  $\mu = p(\text{spike})$  is the overall probability of spiking. The last step in the derivation uses Stein’s lemma, which states that  $\mathbb{E}[\mathbf{x}f(\mathbf{x})] = \mathbb{E}[\nabla f(\mathbf{x})]$  if the expectation is taken over a multivariate Gaussian distribution with identity covariance matrix, corresponding to our white noise stimulus assumption. This calculation thus yields the simple statement that the spike-triggered average is proportional to the gradient (or gain) of the response function, averaged over the input distribution [64]. Applying Stein’s lemma again yields an expression for the STC matrix:

$$\begin{aligned} \mathbf{C}_{\text{STC}} &= \int \mathbf{x}\mathbf{x}^T p(\mathbf{x} \mid \text{spike}) d\mathbf{x} - (\mathbf{x}_{\text{STA}})(\mathbf{x}_{\text{STA}})^T \\ &= \frac{1}{\mu} \int \mathbf{x}\mathbf{x}^T r(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} - (\mathbf{x}_{\text{STA}})(\mathbf{x}_{\text{STA}})^T \\ &= \frac{1}{\mu} \mathbb{E}_{p(\mathbf{x})}[\mathbf{x}\mathbf{x}^T r(\mathbf{x})] - (\mathbf{x}_{\text{STA}})(\mathbf{x}_{\text{STA}})^T \\ &= \frac{1}{\mu} \mathbb{E}_{p(\mathbf{x})}[\nabla^2 r(\mathbf{x})] - (\mathbf{x}_{\text{STA}})(\mathbf{x}_{\text{STA}})^T \end{aligned} \tag{7}$$

Intuitively, these results state that the STA is related to the slope (first derivative) and the STC is related to the Hessian curvature (matrix of second derivatives) of the multi-dimensional nonlinear response function  $r(\mathbf{x})$ .

For example, consider a linear-nonlinear model  $r = f(\mathbf{w}^T \mathbf{x})$  which has the following gradient:  $\nabla r(\mathbf{x}) = f'(\mathbf{w}^T \mathbf{x})\mathbf{w}$  and Hessian:  $\nabla^2 r(\mathbf{x}) = f''(\mathbf{w}^T \mathbf{x})\mathbf{w}\mathbf{w}^T$ . Plugging these expressions into Eqs (6) and (7) reveals that the STA is proportional to  $\mathbf{w}$  and the STC is proportional to  $\mathbf{w}\mathbf{w}^T$ . Therefore, we recover the known result [2] that the STA of the LN model is proportional to the linear filter, and there will be one significant direction in the STC, which is also proportional to the linear filter (with mild assumptions on the nonlinearity,  $f$ , to ensure that slope and curvature terms in (6) and (7) are non-zero).

We can extend this to the case of a multilayered circuit with  $k$  pathways, each of which first filters the stimulus with a filter  $\mathbf{w}_1 \dots \mathbf{w}_k$ . Regardless of how these pathways are then combined, we can write this circuit computation as  $r = f(\mathbf{W}^T \mathbf{x})$  where  $\mathbf{W}$  is a matrix whose columns are the  $k$  pathway filters, and  $f$  is a  $k$ -dimensional time-independent (static) nonlinear function. We can think of the  $k$  dimensional vector  $\mathbf{u} = \mathbf{W}^T \mathbf{x}$  as the activity pattern across each of the  $k$  pathways before any nonlinearity. The gradient for such a model is  $\nabla r(\mathbf{x}) = \mathbf{W}^T \nabla \mathbf{f}(\mathbf{u})$ , where  $\nabla \mathbf{f}(\mathbf{u})$  is the gradient of the  $k$ -dimensional nonlinearity. Using Eq (6), the STA is then a linear combination of the pathway filters:

$$\mathbf{x}_{\text{STA}} = \frac{1}{\mu} \sum_{i=1}^k \alpha_i \mathbf{w}_i,$$

where the weights are given by

$$\alpha_i = \mathbb{E}_{p(\mathbf{x})} [\partial_{u_i} r(\mathbf{u})],$$

and correspond to the average sensitivity, or slope of the neural response  $r$  with respect to changes in the activity of the  $i^{\text{th}}$  filter.

The Hessian for the multilayered model is  $\nabla^2 r(\mathbf{x}) = \mathbf{W} \nabla^2 \mathbf{f} \mathbf{W}^T$ , where  $\nabla^2 \mathbf{f}$  is the  $k$ -by- $k$  matrix of second derivatives of the  $k$ -dimensional nonlinearity  $f(\mathbf{u})$ . From Eq (7), the STC is then given by:

$$\mathbf{C}_{\text{STC}} = \frac{1}{\mu^2} \mathbf{W} \mathbf{H} \mathbf{W}^T, \tag{8}$$

where the  $k$ -by- $k$  matrix  $\mathbf{H}$  is:

$$\mathbf{H} = \mu \mathbb{E}[\nabla^2 f(\mathbf{u})] - \mathbb{E}[\nabla f(\mathbf{u})] \mathbb{E}[\nabla f(\mathbf{u})]^T.$$

This expression implies that nontrivial directions in the column space of  $\mathbf{C}_{\text{STC}}$  correspond to (span the same space as) the column space of  $\mathbf{W}$ . Therefore, the significant eigenvectors of the STC matrix will be linear combinations of the  $k$  pathway filters, and the number of significant eigenvectors is at most  $k$ .

Note that Eqs (6) and (7) are valid for *any* differentiable model, including those with more than two layers, divisive interactions, feedback, and so on.

**Regularized STA.** To compute a regularized STA, without explicitly building an encoding model, we can form an optimization problem that directly denoises the STA:

$$\hat{\mathbf{x}}_{\text{STA}} = \min_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_{\text{STA}}\|_2^2 + \sum_{i=1}^m \gamma_i \phi_i(\mathbf{x}). \tag{9}$$

Here,  $\phi_i(x)$  are the regularization penalty functions, with an associated regularization weight  $\gamma_i$ , and  $\mathbf{x}_{\text{STA}}$  is the raw (sample) STA from recorded data (which is noisy due to finite sampling). Note that we use mean squared error to quantify distance from the raw estimate, but other loss functions may be also used. For the penalty functions  $\phi_i$ , we use an  $\ell_1$  penalty that encourages the estimated filter to be sparse (few non-zero coefficients), and a nuclear norm penalty, which is the sum of the singular values of the spatiotemporal filter  $\mathbf{x}$  when viewed as a spatiotemporal matrix. The nuclear norm penalty is advantageous compared to explicitly forcing the spacetime filter  $\mathbf{x}$  to be low-rank, as it is a “soft” penalty which allows for many small singular values, whereas explicitly forcing the filter to be low-rank forces those to be zero.



**Regularized STC analysis.** The STC eigenvectors are obtained by an eigendecomposition of the STC matrix  $\mathbf{C}$  [10, 76], which is equivalent to solving an optimization problem:

$$\begin{aligned} & \text{maximize} \quad \text{Tr}(\mathbf{U}^T \mathbf{C} \mathbf{U}) \\ & \text{subject to} \quad \mathbf{U}^T \mathbf{U} = \mathbf{I}, \end{aligned} \tag{10}$$

where  $\mathbf{U}$  denotes a matrix whose columns are the orthonormal eigenvectors of  $\mathbf{C}$ . In order to regularize these eigenvectors, we wish to add penalty terms to (10), which precludes a closed form solution to the problem. We circumvent this by reformulating the problem using a convex relaxation. First, we consider the matrix  $\mathbf{X} = \mathbf{U}\mathbf{U}^T$ , corresponding to the outer product of the eigenvectors. Because of the cyclic property of the trace, namely that  $\text{Tr}(\mathbf{U}^T \mathbf{C} \mathbf{U}) = \text{Tr}(\mathbf{U}\mathbf{U}^T \mathbf{C}) = \text{Tr}(\mathbf{X}\mathbf{C})$ , the function to be optimized in (10) depends on the eigenvector matrix  $\mathbf{U}$  only through the combination  $\mathbf{X} = \mathbf{U}\mathbf{U}^T$ . Thus we can directly optimize over the variable  $\mathbf{X}$ . However, the non-convex equality constraint  $\mathbf{U}^T \mathbf{U} = \mathbf{I}$  in (10) is not easily expressible in terms of  $\mathbf{X}$ .  $\mathbf{X}$  is however a projection operator, obeying  $\mathbf{X}^2 = \mathbf{X}$ . We replace this with the constraint that  $\mathbf{X}$  should be contained within the convex hull of the set of rank- $d$  projection matrices. This space of matrices is a convex body known as the *fantope* [77].

The advantage of this formulation is that we obtain a convex optimization problem which can be further augmented with additional functions that penalize the columns of  $\mathbf{X}$  to impose prior knowledge about the structure of the eigenvectors of  $\mathbf{C}$ . Columns of  $\mathbf{X}$  are linear combinations of the eigenvectors of  $\mathbf{C}$ , which are themselves linear combinations of the small set of spatiotemporal filters we are interested in identifying. Therefore, if we expect the spatiotemporal filters of individual biological pathways to have certain structure (for example, smooth, low-rank, or sparse), then we also expect to see those properties in both the eigenvectors and in the columns of  $\mathbf{X}$ .

Putting this logic together, to obtain *regularized* STC eigenvectors, we solve the following convex optimization problem:

$$\hat{\mathbf{X}} = \max_{\mathbf{X}} \text{Tr}(\mathbf{X}\mathbf{C}) + \sum_i \gamma_i \phi_i(\mathbf{X}) \quad \text{subject to} \quad \mathbf{X} \in \mathcal{F}^d. \tag{11}$$

Here  $\mathbf{C}$  is the raw (sample) STC matrix, which is again noisy due to limited recorded data, and  $\mathcal{F}^d$  denotes the fantope, or convex hull of all rank  $d$  projection matrices. Each  $\phi_i$  is a regularization penalty function applied to each of the columns of  $\mathbf{X}$ ; i.e.  $\phi_i(\mathbf{X}) \equiv \sum_{j=1}^N \phi_i(\mathbf{x}^j)$ , where  $\mathbf{x}^j$  denotes the  $j$ 'th column of  $\mathbf{X}$ . Again, we can solve this optimization problem efficiently using proximal consensus algorithms, described above. Common regularization penalties and their corresponding proximal operators are shown in Table 1. The optimization yields a matrix  $\hat{\mathbf{X}}$  in the fantope  $\mathcal{F}^d$ , which may itself have rank higher than  $d$ , so we perform a final eigendecomposition of this matrix to obtain its top-eigenvectors. These eigenvectors constitute our regularized estimate of the eigenvectors of the significant (expansive) STC eigenvectors (to find suppressive directions, one could invert  $\mathbf{C}$  in (11)). A major computational advantage of this formulation is that we only need to store and work with the  $N$  by  $N$  raw STC covariance matrix itself, without ever needing access to the spike-triggered ensemble, an  $N$  by  $M$  matrix where  $M$  (the number of spikes) is typically much greater than  $N$ .

### Subspace overlap

We quantify the overlap between two subspaces as the average of the cosine of the principal (or canonical) angles between the subspaces. The principal angles between two subspaces  $\mathbf{X} \in \mathcal{R}^{n \times p}$  and  $\mathbf{Y} \in \mathcal{R}^{n \times q}$  generalize the idea of angles between vectors. Here we describe a pair

of  $p$  and  $q$  dimensional subspaces in  $n$  dimensional space as the span of the columns of the matrices  $\mathbf{X}$  and  $\mathbf{Y}$ . Assuming without loss of generality that  $p \leq q$ , then we have  $p$  principal angles  $\theta_1, \dots, \theta_p$  that are defined recursively for  $k = 1, \dots, p$  as:

$$\cos \theta_k = \max_{\mathbf{x} \in \mathbf{X}} \max_{\mathbf{y} \in \mathbf{Y}} \mathbf{x}^T \mathbf{y} = \mathbf{x}_k^T \mathbf{y}_k,$$

subject to the constraints that the vectors are unit vectors ( $\mathbf{x}^T \mathbf{x} = \mathbf{y}^T \mathbf{y} = 1$ ) and are orthogonal to the previously identified vectors ( $\mathbf{x}_j^T \mathbf{x} = 0, \mathbf{y}_j^T \mathbf{y} = 0$  for  $j = 1, 2, \dots, k - 1$ ). That is, the first principal angle is found by identifying a unit vector within each subspace such that the correlation, or dot product, between these vectors (these are known as the principal vectors) is maximized. This principal angle is then the inverse cosine of the dot product. Each subsequent principal angle is found by performing the same maximization but restricting each new pair of vectors to be orthogonal to the previous principal vectors in each subspace. The principal angles can be efficiently computed via the QR decomposition [78]. We define subspace overlap as the average of the cosine of the principal angles,  $\frac{1}{p} \sum_{k=1}^p \cos \theta_k$ . This quantity is at most 1 (for two subspaces that span the same space), and at least 0 (for two orthogonal subspaces that share no common directions).

## Supporting information

**S1 Fig. Example bipolar receptive fields.** Each panel shows a spatiotemporal receptive field of an OFF bipolar cell recorded intracellularly from the salamander retina. (TIFF)

**S2 Fig. Subunit and retinal ganglion cell types.** Cell type classification for salamander ganglion cell and subunit filters. (a) K-Means clustering applied to the temporal kernel (temporal component of the spatiotemporal receptive field) of  $n = 23$  recorded retinal ganglion cells. (b) K-Means clustering applied to temporal kernels of  $n = 92$  model-identified subunits. (c) Frequency of the different cell types, both for RGCs and subunits. (TIFF)

## Acknowledgments

The authors would like to thank Ben Naecker, Ben Poole, and Lane McIntosh for discussions as well as Stéphane Deny, Tim Gollisch, Ben Naecker, and Alex Williams for comments on the manuscript.

## Author Contributions

**Conceptualization:** Niru Maheswaranathan, Stephen A. Baccus, Surya Ganguli.

**Data curation:** Niru Maheswaranathan, David B. Kastner.

**Formal analysis:** Niru Maheswaranathan.

**Funding acquisition:** Stephen A. Baccus, Surya Ganguli.

**Investigation:** Niru Maheswaranathan, David B. Kastner, Stephen A. Baccus, Surya Ganguli.

**Methodology:** Niru Maheswaranathan, David B. Kastner, Stephen A. Baccus, Surya Ganguli.

**Resources:** Stephen A. Baccus, Surya Ganguli.

**Software:** Niru Maheswaranathan.

**Supervision:** Stephen A. Baccus, Surya Ganguli.

**Validation:** Niru Maheswaranathan.

**Visualization:** Niru Maheswaranathan.

**Writing – original draft:** Niru Maheswaranathan, Stephen A. Baccus, Surya Ganguli.

**Writing – review & editing:** Niru Maheswaranathan, David B. Kastner, Stephen A. Baccus, Surya Ganguli.

## References

1. Chichilnisky E. A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems*. 2001; 12(2):199–213. <https://doi.org/10.1080/713663221>
2. Paninski L. Convergence properties of three spike-triggered analysis techniques. *Network: Comput Neural Syst*. 2003; 14:437–464. <https://doi.org/10.1088/0954-898X/14/3/304>
3. Schwartz O, Pillow JW, Rust NC, Simoncelli EP. Spike-triggered neural characterization. *Journal of Vision*. 2006; 6(4). <https://doi.org/10.1167/6.4.13> PMID: 16889482
4. Aljadeff J, Lansdell BJ, Fairhall AL, Kleinfeld D. Analysis of neuronal spike trains, deconstructed. *Neuron*. 2016; 91(2):221–259. <https://doi.org/10.1016/j.neuron.2016.05.039> PMID: 27477016
5. Van Steveninck RDR, Bialek W. Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proceedings of the Royal Society of London B: Biological Sciences*. 1988; 234(1277):379–414. <https://doi.org/10.1098/rspb.1988.0055>
6. Agüera y Arcas B, Fairhall AL. What causes a neuron to spike? *Neural Computation*. 2003; 15(8): 1789–1807. <https://doi.org/10.1162/08997660360675044> PMID: 14511513
7. Brenner N, Bialek W, Van Steveninck RdR. Adaptive rescaling maximizes information transmission. *Neuron*. 2000; 26(3):695–702. [https://doi.org/10.1016/S0896-6273\(00\)81205-2](https://doi.org/10.1016/S0896-6273(00)81205-2) PMID: 10896164
8. Schwartz O, Chichilnisky E, Simoncelli EP. Characterizing neural gain control using spike-triggered covariance. *Advances in neural information processing systems*. 2002; 1:269–276.
9. Baccus SA, Meister M. Fast and slow contrast adaptation in retinal circuitry. *Neuron*. 2002; 36(5): 909–919. [https://doi.org/10.1016/S0896-6273\(02\)01050-4](https://doi.org/10.1016/S0896-6273(02)01050-4) PMID: 12467594
10. Fairhall AL, Burlingame AC, Narasimhan R, Harris RA, Puchalla JL, Berry MJ. Selectivity for Multiple Stimulus Features in Retinal Ganglion Cells. *J Neurophysiol*. 2006; 96(5):2724–2738. <https://doi.org/10.1152/jn.00995.2005> PMID: 16914609
11. Rust NC, Schwartz O, Movshon JA, Simoncelli EP. Spatiotemporal Elements of Macaque {V1} Receptive Fields. *Neuron*. 2005; 46(6):945–956. <http://dx.doi.org/10.1016/j.neuron.2005.05.021>. PMID: 15953422
12. Pillow J, Shlens J, Paninski L, Sher A, Litke A, Chichilnisky E, et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008; 454(7207):995–9. <https://doi.org/10.1038/nature07140> PMID: 18650810
13. Heitman A, Brackbill N, Greschner M, Sher A, Litke AM, Chichilnisky E. Testing pseudo-linear models of responses to natural scenes in primate retina. *bioRxiv*. 2016; p. 045336.
14. Latimer KW, Chichilnisky E, Rieke F, Pillow JW. Inferring synaptic conductances from spike trains with a biophysically inspired point process model. In: *Advances in Neural Information Processing Systems*; 2014. p. 954–962.
15. Demb JB, Haarsma L, Freed MA, Sterling P. Functional circuitry of the retinal ganglion cell's nonlinear receptive field. *Journal of Neuroscience*. 1999; 19(22):9756–9767. <https://doi.org/10.1523/JNEUROSCI.19-22-09756.1999> PMID: 10559385
16. Li PH, Field GD, Greschner M, Ahn D, Gunning DE, Mathieson K, et al. Retinal representation of the elementary visual signal. *Neuron*. 2014; 81(1):130–139. <https://doi.org/10.1016/j.neuron.2013.10.043> PMID: 24411737
17. Gollisch T. Features and functions of nonlinear spatial integration by retinal ganglion cells. *Journal of Physiology-Paris*. 2013; 107(5):338–348. <https://doi.org/10.1016/j.jphysparis.2012.12.001>
18. McFarland JM, Cui Y, Butts DA. Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS computational biology*. 2013; 9(7):e1003143. <https://doi.org/10.1371/journal.pcbi.1003143> PMID: 23874185

19. Ozuysal Y, Baccus SA. Linking the computational structure of variance adaptation to biophysical mechanisms. *Neuron*. 2012; 73(5):1002–1015. <https://doi.org/10.1016/j.neuron.2011.12.029> PMID: 22405209
20. Keat J, Reinagel P, Reid RC, Meister M. Predicting Every Spike: A Model for the Responses of Visual Neurons. *Neuron*. 2001; 30(3):803–817. [http://dx.doi.org/10.1016/S0896-6273\(01\)00322-1](http://dx.doi.org/10.1016/S0896-6273(01)00322-1). PMID: 11430813
21. Wu A, Park IM, Pillow JW. Convolutional spike-triggered covariance analysis for neural subunit models. In: *Advances in Neural Information Processing Systems*; 2015. p. 793–801.
22. Freeman J, Field GD, Li PH, Greschner M, Gunning DE, Mathieson K, et al. Mapping nonlinear receptive field structure in primate retina at single cone resolution. *eLife*. 2015; 4:e05241. <https://doi.org/10.7554/eLife.05241> PMID: 26517879
23. Real E, Asari H, Gollisch T, Meister M. Neural Circuit Inference from Function to Structure. *Current Biology*. 2017. <https://doi.org/10.1016/j.cub.2016.11.040> PMID: 28065610
24. Vintch B, Zaharia AD, Movshon JA, Simoncelli EP, et al. Efficient and direct estimation of a neural subunit model for sensory coding. In: *NIPS*; 2012. p. 3113–3121.
25. Park IM, Pillow JW. Bayesian spike-triggered covariance analysis. In: *Advances in neural information processing systems*; 2011. p. 1692–1700.
26. Hochstein S, Shapley RM. Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. *The Journal of Physiology*. 1976; 262(2):265–284. <https://doi.org/10.1113/jphysiol.1976.sp011595> PMID: 994040
27. Victor JD, Shapley RM. The nonlinear pathway of Y ganglion cells in the cat retina. *The Journal of General Physiology*. 1979; 74(6):671–689. <https://doi.org/10.1085/jgp.74.6.671> PMID: 231636
28. Baccus SA. Timing and Computation in Inner Retinal Circuitry. *Annual Review of Physiology*. 2007; 69(1):271–290. <https://doi.org/10.1146/annurev.physiol.69.120205.124451> PMID: 17059359
29. Gollisch T, Meister M. Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina. *Neuron*. 2010; 65(2):150–164. <http://dx.doi.org/10.1016/j.neuron.2009.12.009>. PMID: 20152123
30. Rieke F, Rudd ME. The challenges natural images pose for visual adaptation. *Neuron*. 2009; 64(5):605–616. <https://doi.org/10.1016/j.neuron.2009.11.028> PMID: 20005818
31. Carandini M, Demb JB, Mante V, Tolhurst DJ, Dan Y, Olshausen BA, et al. Do we know what the early visual system does? *The Journal of neuroscience*. 2005; 25(46):10577–10597. <https://doi.org/10.1523/JNEUROSCI.3726-05.2005> PMID: 16291931
32. McIntosh LT, Maheswaranathan N, Nayebi A, Ganguli S, Baccus SA. Deep learning models of the retinal response to natural scenes. In: *Advances in Neural Information Processing Systems*; 2016. p. 1361–1369.
33. Maheswaranathan N, McIntosh LT, Kastner DB, Melander J, Brezovec L, Nayebi A, Wang J, Ganguli S, Baccus SA. Deep learning models reveal internal structure and diverse computations in the retina under natural scenes. *bioRxiv*. 2018; p. 340943.
34. Demb JB, Zaghoul K, Haarsma L, Sterling P. Bipolar cells contribute to nonlinear spatial summation in the brisk-transient (Y) ganglion cell in mammalian retina. *The Journal of neuroscience*. 2001; 21(19):7447–7454. <https://doi.org/10.1523/JNEUROSCI.21-19-07447.2001> PMID: 11567034
35. Turner MH, Rieke F. Synaptic rectification controls nonlinear spatial integration of natural visual inputs. *Neuron*. 2016; 90(6):1257–1271. <https://doi.org/10.1016/j.neuron.2016.05.006> PMID: 27263968
36. Werblin FS. Six different roles for crossover inhibition in the retina: correcting the nonlinearities of synaptic transmission. *Visual neuroscience*. 2010; 27(1-2):1–8. <https://doi.org/10.1017/S0952523810000076> PMID: 20392301
37. Enroth-Cugell C, Robson JG. The contrast sensitivity of retinal ganglion cells of the cat. *The Journal of physiology*. 1966; 187(3):517–552. <https://doi.org/10.1113/jphysiol.1966.sp008107> PMID: 16783910
38. Ölveczky BP, Baccus SA, Meister M. Segregation of object and background motion in the retina. *Nature*. 2003; 423(6938):401–408. <https://doi.org/10.1038/nature01652> PMID: 12754524
39. Geffen MN, de Vries SEJ, Meister M. Retinal Ganglion Cells Can Rapidly Change Polarity from Off to On. *PLoS Biol*. 2007; 5(3):e65. <https://doi.org/10.1371/journal.pbio.0050065> PMID: 17341132
40. Pillow JW, Simoncelli EP. Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis. *Journal of vision*. 2006; 6(4):9–9. <https://doi.org/10.1167/6.4.9>
41. Aljadeff J, Segev R, Berry MJ II, Sharpee TO. Spike triggered covariance in strongly correlated Gaussian stimuli. *PLoS Comput Biol*. 2013; 9(9):e1003206. <https://doi.org/10.1371/journal.pcbi.1003206> PMID: 24039563

42. Sharpee TO. Computational identification of receptive fields. *Annual review of neuroscience*. 2013; 36:103. <https://doi.org/10.1146/annurev-neuro-062012-170253> PMID: 23841838
43. Liu YS, Stevens CF, Sharpee TO. Predictable irregularities in retinal receptive fields. *Proceedings of the National Academy of Sciences*. 2009; 106(38):16499–16504. <https://doi.org/10.1073/pnas.0908926106>
44. Kaardal J, Fitzgerald JD, Berry MJ, Sharpee TO. Identifying functional bases for multidimensional neural computations. *Neural computation*. 2013; 25(7):1870–1890. [https://doi.org/10.1162/NECO\\_a\\_00465](https://doi.org/10.1162/NECO_a_00465) PMID: 23607565
45. Bölinger D, Gollisch T. Closed-Loop Measurements of Iso-Response Stimuli Reveal Dynamic Nonlinear Stimulus Integration in the Retina. *Neuron*. 2012; 73(2):333–346. <http://dx.doi.org/10.1016/j.neuron.2011.10.039>. PMID: 22284187
46. Liu JK, Schreyer HM, Onken A, Rozenblit F, Khani MH, Krishnamoorthy V, et al. Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nature Communications*. 2017; 8(1):149. <https://doi.org/10.1038/s41467-017-00156-9> PMID: 28747662
47. Park M, Pillow JW. Bayesian inference for low rank spatiotemporal neural receptive fields. In: *Advances in Neural Information Processing Systems*; 2013. p. 2688–2696.
48. Thorson IL, Liénard J, David SV. The essential complexity of auditory receptive fields. *PLoS computational biology*. 2015; 11(12):e1004628. <https://doi.org/10.1371/journal.pcbi.1004628> PMID: 26683490
49. Schwartz GW, Okawa H, Dunn FA, Morgan JL, Kerschensteiner D, Wong RO, et al. The spatial structure of a nonlinear receptive field. *Nature neuroscience*. 2012; 15(11):1572. <https://doi.org/10.1038/nn.3225> PMID: 23001060
50. Asari H, Meister M. Divergence of visual channels in the inner retina. *Nature neuroscience*. 2012; 15(11):1581–1589. <https://doi.org/10.1038/nn.3241> PMID: 23086336
51. Asari H, Meister M. The projective field of retinal bipolar cells and its modulation by visual context. *Neuron*. 2014; 81(3):641–652. <https://doi.org/10.1016/j.neuron.2013.11.029> PMID: 24507195
52. Rudelson M, Vershynin R. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM (JACM)*. 2007; 54(4):21. <https://doi.org/10.1145/1255443.1255449>
53. Wu SM, Gao F, Maple BR. Functional architecture of synapses in the inner retina: segregation of visual signals by stratification of bipolar cell axon terminals. *The Journal of Neuroscience*. 2000; 20(12):4462–4470. <https://doi.org/10.1523/JNEUROSCI.20-12-04462.2000> PMID: 10844015
54. Wässle H, Puller C, Müller F, Haverkamp S. Cone contacts, mosaics, and territories of bipolar cells in the mouse retina. *The Journal of Neuroscience*. 2009; 29(1):106–117. <https://doi.org/10.1523/JNEUROSCI.4442-08.2009> PMID: 19129389
55. Jadzinsky PD, Baccus SA. Synchronized amplification of local information transmission by peripheral retinal input. *eLife*. 2015; 4. <https://doi.org/10.7554/eLife.09266> PMID: 26568312
56. Attneave F. Some informational aspects of visual perception. *Psychological review*. 1954; 61(3):183. <https://doi.org/10.1037/h0054663> PMID: 13167245
57. Barlow HB. In: *Possible principles underlying the transformations of sensory messages*. Cambridge, MA: MIT Press; 1961. p. 217–234.
58. Hyvärinen A, Hurri J, Hoyer PO. *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*. vol. 39. Springer Science & Business Media; 2009.
59. Atick JJ, Redlich AN. Towards a theory of early visual processing. *Neural Computation*. 1990; 2(3):308–320. <https://doi.org/10.1162/neco.1990.2.3.308>
60. Franke K, Berens P, Schubert T, Bethge M, Euler T, Baden T. Inhibition decorrelates visual feature representations in the inner retina. *Nature*. 2017; 542(7642):439–444. <https://doi.org/10.1038/nature21394> PMID: 28178238
61. Pitkow X, Meister M. Decorrelation and efficient coding by retinal ganglion cells. *Nature neuroscience*. 2012; 15(4):628–635. <https://doi.org/10.1038/nn.3064> PMID: 22406548
62. Berry MJ, Warland DK, Meister M. The structure and precision of retinal spike trains. *Proceedings of the National Academy of Sciences*. 1997; 94(10):5411–5416. <https://doi.org/10.1073/pnas.94.10.5411>
63. Sterling P, Laughlin S. *Principles of neural design*. MIT Press; 2015.
64. Bialek W, van Steveninck RR. Features and dimensions: Motion estimation in fly vision. *arXiv preprint q-bio/0505003*. 2005.
65. Atick JJ, Redlich AN. What does the retina know about natural scenes? *Neural computation*. 1992; 4(2):196–210. <https://doi.org/10.1162/neco.1992.4.2.196>
66. Field GD, Rieke F. Nonlinear signal transfer from mouse rods to bipolar cells and implications for visual sensitivity. *Neuron*. 2002; 34(5):773–785. [https://doi.org/10.1016/S0896-6273\(02\)00700-6](https://doi.org/10.1016/S0896-6273(02)00700-6) PMID: 12062023

67. Bialek W. Biophysics: searching for principles. Princeton University Press; 2012.
68. Clark DA, Benichou R, Meister M, da Silveira RA. Dynamical adaptation in photoreceptors. *PLOS Comput Biol*. 2013; 9(11):e1003289. <https://doi.org/10.1371/journal.pcbi.1003289> PMID: 24244119
69. Kim KJ, Rieke F. Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. *The Journal of Neuroscience*. 2001; 21(1):287–299. <https://doi.org/10.1523/JNEUROSCI.21-01-00287.2001> PMID: 11150346
70. Manookin MB, Demb JB. Presynaptic mechanism for slow contrast adaptation in mammalian retinal ganglion cells. *Neuron*. 2006; 50(3):453–464. <https://doi.org/10.1016/j.neuron.2006.03.039> PMID: 16675399
71. Weick M, Demb JB. Delayed-rectifier K channels contribute to contrast adaptation in mammalian retinal ganglion cells. *Neuron*. 2011; 71(1):166–179. <https://doi.org/10.1016/j.neuron.2011.04.033> PMID: 21745646
72. Kastner DB, Baccus SA. Spatial segregation of adaptation and predictive sensitization in retinal ganglion cells. *Neuron*. 2013; 79(3):541–554. <https://doi.org/10.1016/j.neuron.2013.06.011> PMID: 23932000
73. Naecker B, Maheswaranathan N, Ganguli S, Baccus SA. Pyret: A Python package for analysis of neurophysiology data. *The Journal of Open Source Software*. 2017. <https://doi.org/10.21105/joss.00137>
74. Parikh N, Boyd S. Proximal algorithms. *Foundations and Trends in optimization*. 2013; 1(3):123–231.
75. Polson NG, Scott JG, Willard BT, et al. Proximal Algorithms in Statistics and Machine Learning. *Statistical Science*. 2015; 30(4):559–581. <https://doi.org/10.1214/15-STS530>
76. Rajan K, Marre O, Tkačik G. Learning quadratic receptive fields from neural responses to natural stimuli. *Neural computation*. 2013; 25(7):1661–1692. [https://doi.org/10.1162/NECO\\_a\\_00463](https://doi.org/10.1162/NECO_a_00463) PMID: 23607557
77. Vu VQ, Cho J, Lei J, Rohe K. Fantope Projection and Selection: A near-optimal convex relaxation of sparse PCA. In: *Advances in Neural Information Processing Systems 26*; 2013. p. 2670–2678.
78. Björck Å, Golub GH. Numerical methods for computing angles between linear subspaces. *Mathematics of computation*. 1973; 27(123):579–594. <https://doi.org/10.1090/S0025-5718-1973-0348991-3>