



# Stimulus-invariant aspects of the retinal code drive discriminability of natural scenes

Benjamin D. Hoshaj<sup>a,1</sup>, Caroline M. Holmes<sup>b,1,2</sup>, Kyle Bojanek<sup>a</sup>, Jared M. Salisbury<sup>c,d</sup>, Michael J. Berry II<sup>e</sup>, Olivier Marre<sup>f</sup>, and Stephanie E. Palmer<sup>a,c,d,g,3</sup>

Affiliations are included on p. 9.

Edited by J. Anthony Movshon, New York University, New York, NY; received August 8, 2023; accepted November 11, 2024

Everything that the brain sees must first be encoded by the retina, which maintains a reliable representation of the visual world in many different, complex natural scenes while also adapting to stimulus changes. This study quantifies whether and how the brain selectively encodes stimulus features about scene identity in complex naturalistic environments. While a wealth of previous work has dug into the static and dynamic features of the population code in retinal ganglion cells (RGCs), less is known about how populations form both flexible and reliable encoding in natural moving scenes. We record from the larval salamander retina responding to five different natural movies, over many repeats, and use these data to characterize the population code in terms of single-cell fluctuations in rate and pairwise couplings between cells. Decomposing the population code into independent and cell-cell interactions reveals how broad scene structure is encoded in the retinal output, while the single-cell activity adapts to different stimuli, the population structure captured in the sparse, strong couplings is consistent across natural movies as well as synthetic stimuli. We show that these interactions contribute to encoding scene identity. We also demonstrate that this structure likely arises in part from shared bipolar cell input as well as from gap junctions between RGCs and amacrine cells.

neural coding | natural scenes | retina | neural computation | population code

While single cells individually encode specific stimulus features (1–3), it is their aggregate response that drives our perception (4–7). For this reason, it is important to understand not only how individual cells respond to stimuli but also how cells influence each other within a population (8–11). Significant theoretical work has been devoted to understanding population responses (12–16), in tandem with experimental innovations in recording from a large number of cells simultaneously (17–20). This work has shown that relatively small, diffuse neural correlations are crucial for reproducing the distributions of observed population states or “vocabulary” of the brain. Open questions remain about the functional consequence of these correlations for an organism behaving in its natural environment. This foundation in both theory and experiment on neural population codes creates an opportunity to move to more complex, dynamic stimuli and analyze the code in terms of the readout goals of downstream networks.

The natural environment has many complex spatiotemporal features that make neural encoding in the wild difficult to quantify and assess. Natural scenes vary in luminance over many orders of magnitude (21) and variance (22, 23) and have complicated temporal and spatial structure (24, 25). Visual systems adapt to these changes on many scales in time and space. Neural systems show near-perfect adaptation to these changes (26), so a question remains about how brains recover scene-specific information once in an adapted state. This open question has led many studies to investigate animal behavior in natural settings (27–30). In this work, we quantify the structure of the neural code at the input end, and how it might support downstream readout that ultimately drives behavior in complex environments.

We use the spike activity in a salamander retina in response to natural stimuli to infer a minimal model of the population structure of its output retinal ganglion cells (RGCs). This structure is, as measured both by statistical predictions and by functional consequences, sparse and conserved across scenes. A human observer of these natural scenes can immediately tell that they are different: how does a small patch of the retina capture scene identity and convey it to the brain? We show that the sparse, consistent correlations between cells helps the population carry long time-scale information about broad scene statistics. Finally, we show that these sparse conserved couplings appear to

## Significance

Visual perception relies on the reliable representation of the world and starts in the retina. While single cells encode specific, local stimulus features, it is the activity of the population as a whole that drives perception and action. While previous work has shown that populations have widespread correlations, the functional consequence of these interactions is less understood. Here we separate the population retinal code into independent and cell-cell interactions to reveal how underlying population structure improves encoding in natural dynamic scenes. We demonstrate that population level interactions are crucial for downstream readout of scene identity. These cell-cell interactions arise directly from the retinal anatomy itself, suggesting a mechanistic underpinning for population level encoding of visual scenes.

Author contributions: B.D.H., C.M.H., K.B., M.J.B., O.M., and S.E.P. designed research; O.M. and S.E.P. performed research; B.D.H., C.M.H., and K.B. contributed new reagents/analytic tools; B.D.H., C.M.H., K.B., J.M.S., and S.E.P. analyzed data; and B.D.H., C.M.H., K.B., and S.E.P. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2024 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<sup>1</sup>B.D.H. and C.M.H. contributed equally to this work.

<sup>2</sup>Present address: Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138.

<sup>3</sup>To whom correspondence may be addressed. Email: [sepalmer@uchicago.edu](mailto:sepalmer@uchicago.edu).

Published December 19, 2024.

arise from both shared input (bipolar and amacrine cells) and direct connections (gap junctions).

## Results

### Probing Multiple Naturalistic, Dynamic Inputs to the Retina.

We took dense extracellular recordings from RGCs in the larval tiger salamander (Fig. 1A) while presenting 20-s clips from five different movies (Fig. 1B), in a random order, with each shown more than 80 times (Fig. 1C). Salamanders undergo metamorphosis, exposing them to both underwater and terrestrial environments while their retinal structure remains largely the same (31, 32), and navigate through their environment, generating self-motion. The movies chosen represent a sampling of the wide variety of scenes that occur in the organism's ecological niche (Fig. 1B).

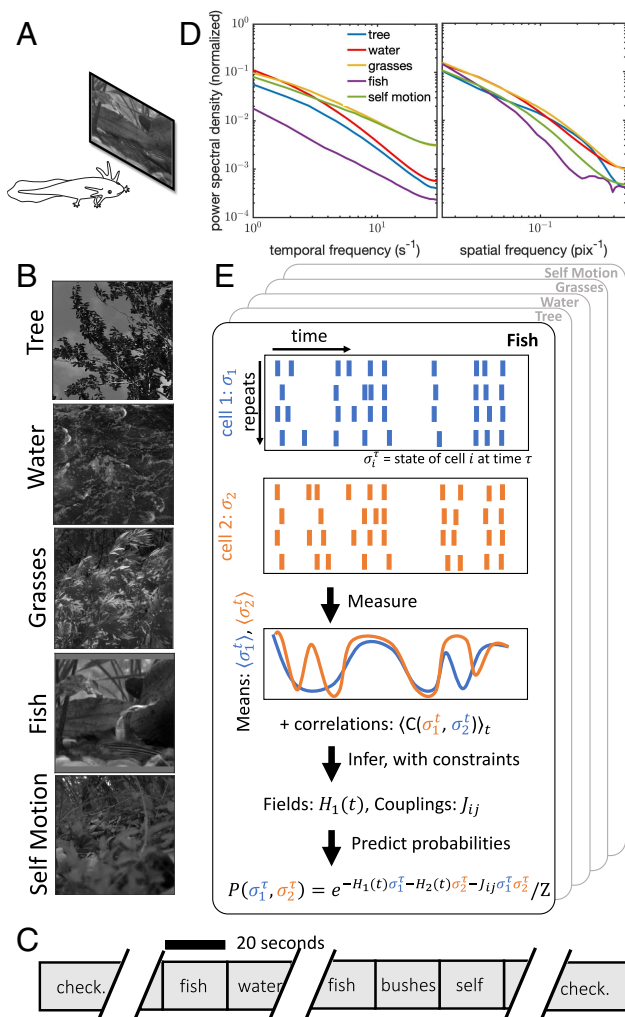
Because of the diversity of subjects and locations, the movies exhibit significantly different temporal and spatial frequency spectra (Fig. 1D). As luminance correlations alone fail to capture behaviorally relevant features like motion, work has quantified the substantially different shapes and timescales of velocity autocorrelations for different scenes (33, 34). As animals navigate through different environments, they rapidly adapt to these changes in stimulus statistics (26, 35, 36). Thus, playing a variety of natural scenes in a single experimental session probes a fuller range of neural response repertoires and gives us an opportunity to capture the true range of the brain's code.

**Pairwise Couplings Are Consistent Across Movies.** To fully describe a dynamic binary population code, we must enumerate  $2^N$  possible states at each time in the response. Even for modest  $N$ , a fully expressive population code is both experimentally inaccessible and potentially unreadable by downstream networks. To summarize the population code, we use maximum entropy modeling, which has a long history of success using  $O(N^2)$  parameters to capture the structure in biological data, even recapitulating higher-order features not explicitly constrained by the model (37–44). We chose these “maxent” models over a generalized linear model (GLM) approach. While GLMs have led to significant insight into adaptation filters in individual cells and effective interactions between cells (10, 45–47), they are not well suited for analyzing the structure of the retinal population code in natural scenes, particularly at the time resolution of interest (16.67 ms, set by the movie frame rate). At this temporal resolution, the primary correlation between cells is the zero-lag correlation, which cannot be modeled by GLMs. In addition, as generative models, GLMs regularly struggle with self-excitation (48).

In applications of maximum entropy techniques in neuroscience, these models are typically constrained by the average firing rate of each cell and the correlation between each pair of cells (37, 49),  $\langle \sigma_i \rangle$  and  $\langle \sigma_i \sigma_j \rangle$ . We use a time-dependent maximum entropy model (50) that is constrained by the time-varying firing rates averaged across repeated stimuli. Our model takes the form

$$P(\vec{\sigma}^t) = \frac{1}{Z} e^{-\sum_i h_i^t \sigma_i^t - \sum_{i < j} J_{ij} \sigma_i^t \sigma_j^t}, \quad [1]$$

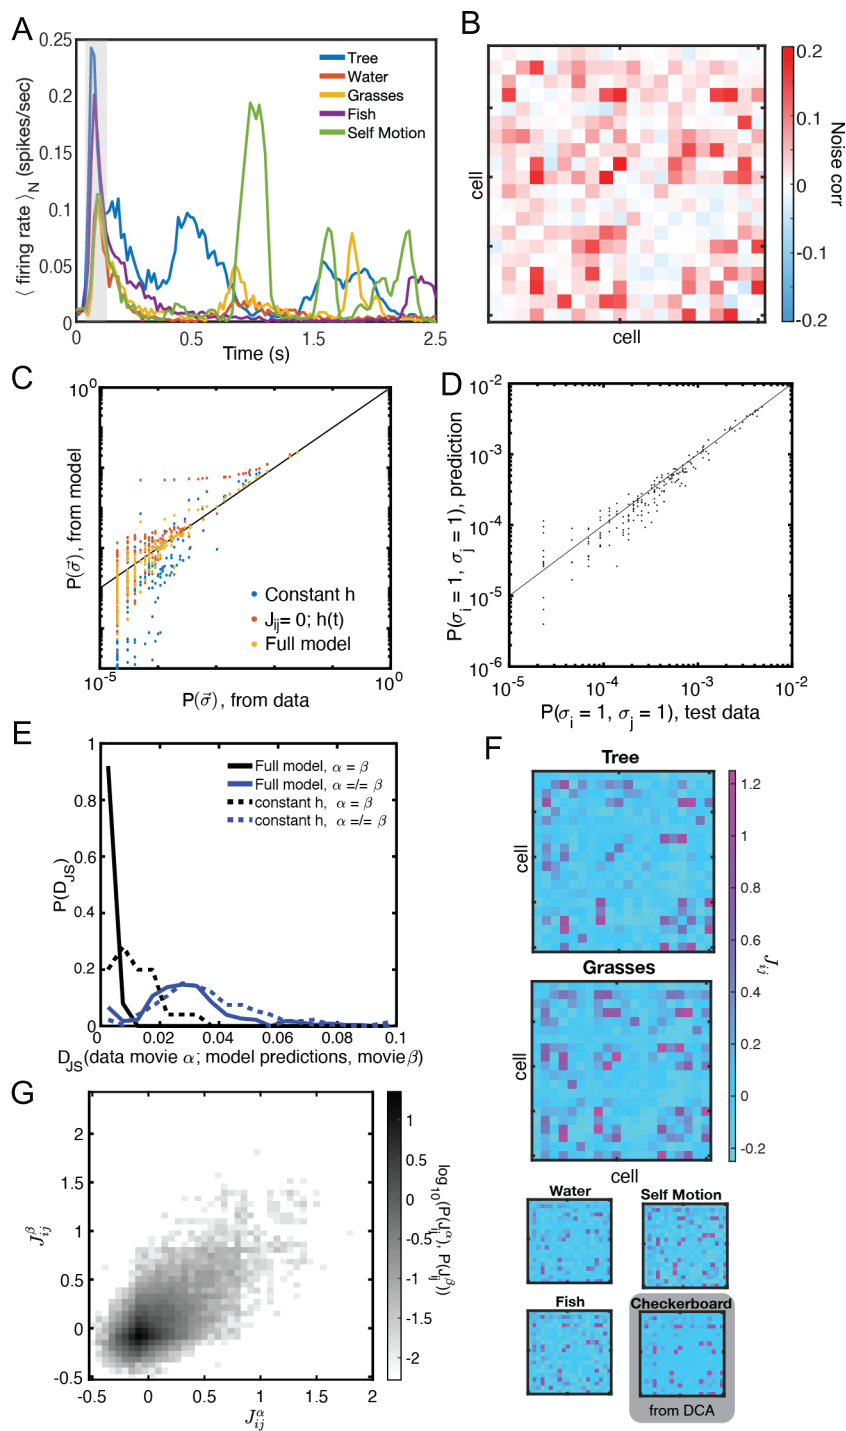
and our constraints are on  $\langle \sigma_i^t \rangle_k$ , which capture each cell's individual response to the stimulus at time  $t$  averaged over trials,  $k$ , (Fig. 2A) as well as  $\langle \sigma_i^t \sigma_j^t \rangle_{t,k}$ , the correlations between cells (Fig. 2B). These two constraints map to two sets of parameters,



**Fig. 1.** Measuring RGC responses to diverse natural scenes. (A) Voltage responses were recorded from the RGC layer of a salamander retina stimulated by natural movies. (B) Example frames from each of five natural scenes, which show, respectively, trees blowing in the wind; flowing water; ferns and grasses waving back and forth in a breeze; fish swimming; and woodland underbrush as viewed by a moving camera. (C) In order to probe the statistics of responses, natural scenes were repeated a minimum of 80 times in pseudorandom order. A checkerboard stimulus lasting 30 min was shown before and after the natural scenes. (D) The five natural movies used have different statistics, including (shown here) spatial and temporal frequency spectra. (E) To model responses to each of the movies, time-dependent maximum entropy models are fit to each of the five natural scenes.

the time-dependent fields  $h_i^t$  and the static couplings  $J_{ij}$ . Interactions between the time-dependent fields  $h_i^t$  absorb any stimulus-dependent correlations, leaving the couplings  $J_{ij}$  to capture the noise correlations between cells. We independently train models for each movie, fixing couplings as constant across time within movies but allowing them to change across movies (Fig. 1E). This means that couplings are indexed by cell and movie ( $J_{ij}^\alpha$ ), while fields are indexed by cell, movie, and time ( $h_i^{t,\alpha}$ ).

Our models accurately predict the full suite of population state probabilities (Fig. 2C), which indicates that higher-order information in the data is well accounted for by these second-order constraints. We can also use more traditional train-and-test statistics to validate our model and find that when it is trained on half of the data, it is able to predict the pairwise statistics of the held-out half (Fig. 2D). As in other pairwise



**Fig. 2.** RGC populations maintain consistent couplings across a variety of natural stimuli. (A) Average population firing rate as a function of time for the first 2 s of each natural movie. In the first 200 ms (gray bar), the population firing rate dramatically peaks; after that, rates remain variable in response to stimulus statistics. (B) In addition to time-varying firing rates, cells are characterized by their noise correlations, here shown for a single 20-cell group of cells responding to the tree movie. (C) Probability of population states for a single group of 20 cells, as measured from data and compared to models. The full model successfully extrapolates from second-order statistics to population state. Models with either no couplings or with constant fields extrapolate less well. (D) Prediction of two-neuron statistics on data held out of the training procedure. Our model, when fit on half of the data, is able to successfully predict the pairwise statistics of the remaining half of the data. (E) Across a set of ten random choices of 20-cell groups and all movies, our model predicts population states well (low  $D_{JS}$ , black line) and the traditional constant-field model does less well (dotted black line). Models across movies are substantially different from each other and do not predict states well between movies (blue lines). (F) Couplings across all five natural scenes and the checkerboard white noise stimulus, for a single choice of group. In all cases, couplings are consistent across stimuli, especially for the strongest couplings. Even for the checkerboard, where another modeling procedure [direct coupling analysis (DCA)] was used to fit the interactions, similar couplings are found. (G) Couplings are quantitatively similar across movies ( $R^2 = 0.74$ ), especially for the significantly nonzero couplings. Here, we show the values of the couplings  $J_{ij}$  for all pairs of movies, for ten different groups of 20 cells each. Instead of showing those many thousands of individual points, we have binned the  $J$ s and displayed the resulting values as a 2D probability density, on a log scale (gray scale bar). The SD on these  $J$  parameters, estimated through bootstrapping, is roughly 0.11.

models fit to neural data, a relatively small number of parameters can capture the full response repertoire of the neural population (37). This fit is significantly better than the traditional constant-field models without a time-varying field. To quantify this fit, we calculate the overall divergence between the model probabilities and the observed state frequencies,  $D_{JS}(P_{model}, P_{data})$  (Fig. 2E). This analysis again shows that our full model matches our data, while models with constant fields do significantly worse.

Given the quality of predictions from our models, and that the fields are explicitly constrained by stimulus-induced single-cell statistics, we assume that the interaction matrices  $J_{ij}$  carry the

essence of the intrinsic population structure, while the fields carry information about independent responses to the external visual stimulus. We find that these  $J_{ij}$  matrices appear to have similar structure across movies (Fig. 2F). The consistent couplings are not unique to the particular 20-cell group analyzed in Fig. 2F. For a selection of randomly chosen groups, we plot the coupling  $J_{ij}^\alpha$  between cells  $i$  and  $j$  in movie  $\alpha$  against the couplings  $J_{ij}^\beta$ . We observe a strong correlation between the sparse set of large-magnitude cell–cell interactions from movie to movie (Fig. 2G). These can be compared to the SD of our estimate of  $J_{ij}$  in the tree movie, which, as determined by bootstrapping are largely independent of  $J_{ij}$  and have an average value of 0.11.

Previous work has similarly found consistent cell–cell couplings or cell–cell noise correlations across visual inputs (50, 52, 53). Our goal is to dissect this sparse, consistent “skeleton” of interactions more fully and explore both the origins and functional consequences of this structure. In this study, the correlation between couplings has been demonstrated across a range of naturalistic, dynamic stimuli. This consistency is particularly surprising given that these are independently trained models, which separately learned similar couplings despite substantially different scene statistics (Fig. 1D) and predicted population activity (Fig. 2E). This implies that this structure arises from the retina itself, rather than being inherited from the stimulus.

To test whether the similar couplings might arise due to shared long-timescale and length-scale correlations in natural scenes (24), or if they are a property of the feedforward, functional anatomy of the retina, we compare the couplings fit in the natural movies to those found in response to an entirely different, nonnaturalistic stimulus.

The population response to a white-noise checkerboard stimulus also has the same kind of pairwise interactions we inferred in response to the natural movies (Fig. 2F). Due to a lack of repeats of the white noise stimulus, we used a different method, DCA (54), to infer these couplings (see *Materials and Methods* for details). Despite the many differences between this model and those fit to the natural movies (both stimulus and model details), we extract the same strong coupling structure. This strengthens the argument that the observed couplings are indicative of real biological interactions, not correlations inherited from the input.

**Sparse, Consistent Couplings Capture Population Response Structure.** The distribution of values in the  $J_{ij}$  matrices is heavy-tailed (Fig. 3A, *Inset*), which means that there are rare, sparse strong values. This sparsity of correlated, strong couplings connects with previous work that has highlighted the possible functional role of heavy-tailed coupling distributions in neural networks (55). Ideas from sparse coding theory have permeated discussions of how large neural populations both store and retrieve information (56–58). We will quantify this sparsity and consistency, and investigate their role in determining the retinal population code.

One measure of the effective sparsity of the  $J_{ij}$  matrices is to check the role of the small set of large couplings in determining population states. Specifically, we set 90% of the smaller magnitude couplings (both positive and negative) to zero. We then refit the remaining, strongest 10% of the  $J_{ij}$ s (an example is shown in Fig. 3A). The resulting skeleton model must deviate from the full model in its fit to the true population response probabilities, but it does significantly better than a conditionally independent model with no  $J$ 's at all (Fig. 3B for an example of one 20 cell group).

We can broaden this analysis to look at the  $D_{JS}$  between the data and our model-predicted probability distributions, averaged over the choice of 20-cell group (Fig. 3C). We find that the  $D_{JS}$  is small between the data and the full model (as previously shown in Fig. 2E), and increases significantly for a model without couplings. The  $D_{JS}$  for the skeleton model is nearly identical to that of the full model, meaning that this sparse set of couplings captures almost all of the population response structure embedded in the full  $J$  matrix.

We can use this analysis to investigate the role of the *changes* in the couplings on the predicted response distributions. We do this by making chimera models, where  $h_i^{t,\alpha}$ 's are fit as in the full

model, but couplings are inherited from a different movie,  $J_{ij}^\beta$  (Fig. 3D). If the model with fields  $h_i^{t,\alpha}$  and  $J_{ij}^{\beta \neq \alpha}$  makes very different predictions from a model with  $h_i^{t,\alpha}$  and  $J_{ij}^\alpha$ , it means that the  $J$ 's specific to that movie matter. Instead, we find that models with the same  $\alpha$  and various  $\beta$  choices perform similarly (evidenced by the horizontal bands in the right part of Fig. 3E). The differences in coupling have minimal effects on determining population activity. Importantly, this is not a statement that couplings do not matter—instead, the *variation* in the couplings has minimal effect, and that they seem to be consistent across scenes.

**Sparse, Consistent Couplings Allow for Faster Decoding of Scene Identity.**

While the results in Fig. 2 show that couplings are consistent across movies, it is not clear from that analysis alone what this means for downstream readout. In general, while couplings between cells have been inferred and their structure discussed in a wide range of neural populations (37, 51, 59, 60), the functional consequence of these couplings is less explored. The couplings observed in these data are both consistent across natural scenes and heavy-tailed (Fig. 3A). Previous work has shown that weak couplings combine to have a large effect on population activity (37), while others have suggested that ignoring weak interactions has minimal effect on population responses (40, 42). Exploring a broad set of visual scenes allows us to tease apart these two seemingly competing claims about the functional consequence of cell–cell couplings in the retina.

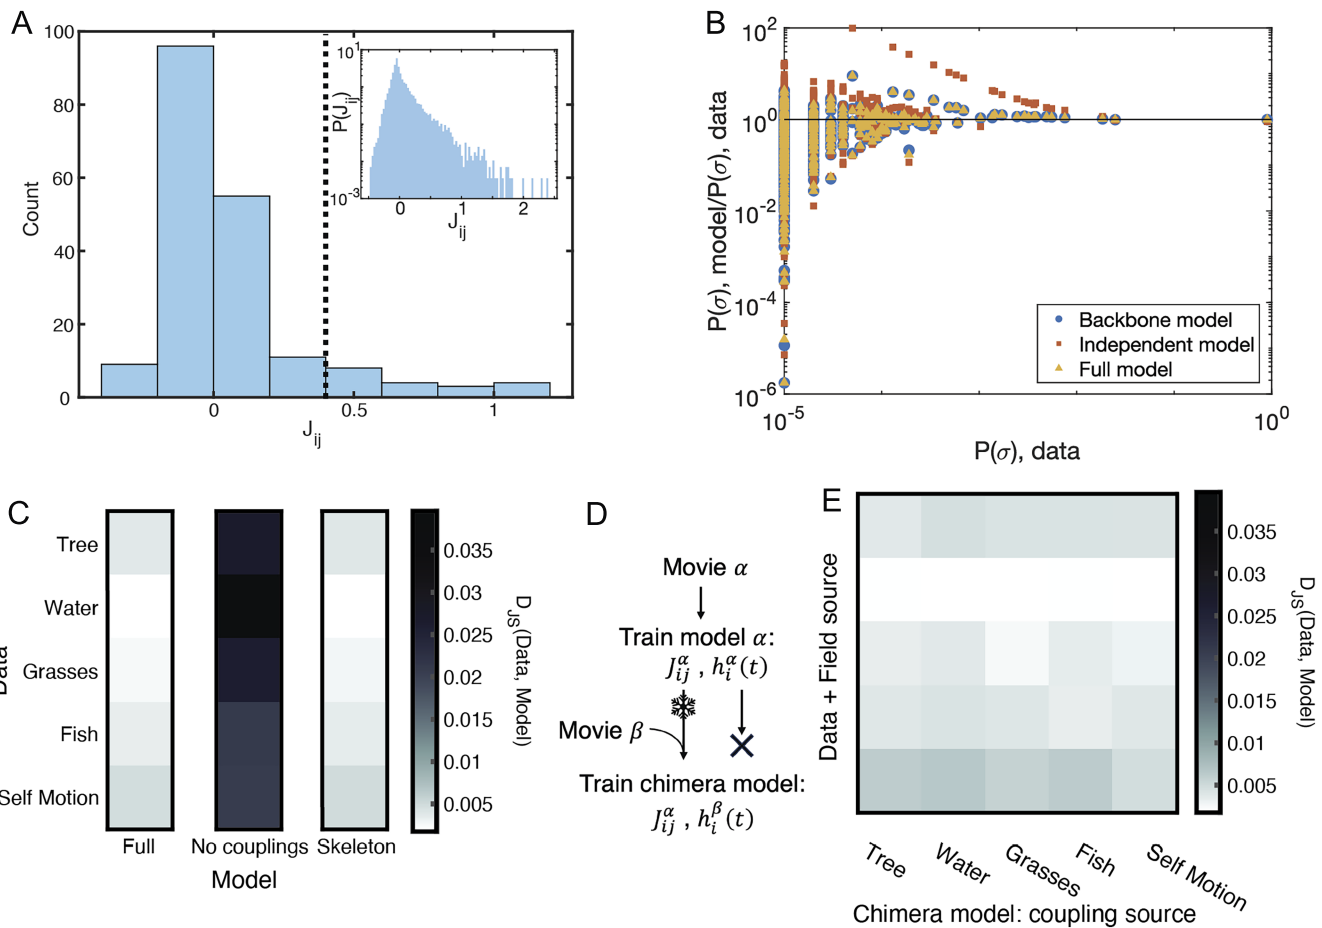
While our preceding analysis shows that the core predictive power of our coupling models is invariant to elimination of all but the largest 10% of interactions, and that the changes in coupling observed across movies do not significantly perturb model predictions, the real biological system might *not* care about inferring the population state distribution. A more compelling test of the role of couplings would be to investigate their role in a downstream behavioral or cognitive task.

Previous work has investigated the effect of interactions between retinal cells on stimulus encoding and retinal development (8, 12, 44, 61–63). In contrast with this work, we focus on the decoding problem of inferring scene identity. The sparse, strong couplings may combine with the time-dependent fields to signal broad scene spatiotemporal structure. The movies contain a suite of higher-order features that make each one readily discriminable to the human eye but may also impact the local, correlated retinal population code in a readable way.

To test whether the couplings affect the discrimination of scenes, we take advantage of the fact that each of our movies comes from a significantly different environment, which means that the information about scene statistics can be approximated by information about scene identity (Fig. 1D). We can then quantify the accumulation of evidence about scene identity in a Bayesian framework, given an ideal decoder with the posterior  $P(\text{movie}|\text{spikes})$ . This decoding task is similar to the real problem solved by downstream brain areas when an organism moves between scenes as it navigates its natural environment. The brain must invoke different behavioral repertoires and priors in different scenes.

We use a Bayesian approach to measure  $P(\text{movie}|\text{spikes})$ ,

$$P(\text{movie}|\text{spikes}) = \frac{P(\text{spikes}|\text{movie})P(\text{movie})}{P(\text{spikes})}. \quad [2]$$



**Fig. 3.** The conserved, sparse coupling structure captures population activity. (A) The distribution of  $J_{ij}$ s is heavy-tailed (*Inset*), and in any particular 20-cell group there are a small number of large couplings. We study this by choosing the top 10% of couplings (above the dotted line), which we will refer to as the skeleton. In general, these couplings can be either positive or negative. (B) Probabilities of population states, for our full model, an independent model (all  $J_{ij} = 0$ ), and a model trained with only the skeleton values allowed to be nonzero. The full and skeleton models perform nearly identically. (C) Across movies and 10 different groups of 20 cells, the skeleton model performs as well as the full model (i.e., the left column is similar to the right column). (D) To investigate the effects of the variation in couplings, we define chimera models, where the couplings are learned from one movie, and then the fields are trained from a second movie. (E) The chimera models perform nearly as well as the full model (each column is nearly identical to the full model column in panel C). We note that the diagonal entries of this matrix are exactly the values in the full model column of panel C.

In order to calculate this quantity, we sample population activity from each model that we fit, either making sparse  $J$ 's, eliminating couplings entirely, or mixing couplings between scenes (see *Materials and Methods* for details). Importantly, because we generate data from reduced or mixture models at will, we can quantify how much each component of the model matters for inference of scene identity.

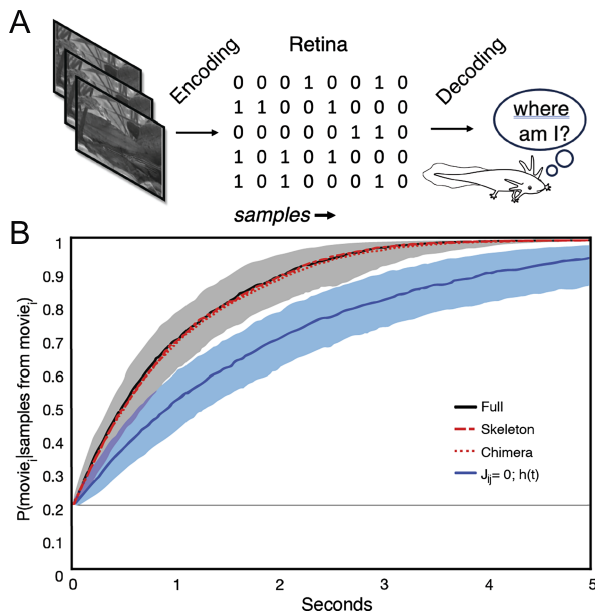
We find that samples generated from even a 20-cell group carry enough information to correctly identify a scene within a few seconds (Fig. 4B, solid black line). With more neurons, even faster decoding is likely possible. By contrast, spikes generated from a conditionally independent model, with all  $J$ 's set to zero, take nearly twice as long to achieve the same confidence level in scene identity. We emphasize that the only difference between these models is the couplings; both models have time-varying fields, which capture the spatial and temporal correlations in the natural scene. The disparity in performance is surprising, as it was not a priori obvious that couplings would or should contribute to decoding (64). Prior efforts to establish the functional role of couplings, especially in natural scenes, have shown that correlations matter for discriminating between spontaneous and stimulus-driven responses (14). To push this further, we can

examine how couplings support discrimination *between* complex stimulus-driven ensembles.

A sparse coupling matrix also supports fast inference of scene identity. We calculate the accumulation of evidence from samples from a skeleton model, where only 10% of couplings are nonzero. Evidence about scene identity accumulates just as rapidly as that of the full model (Fig. 4B, dot-dash red line). The benefit to decoding that arises from the couplings is solely contained in the few strongest couplings, and weaker couplings do not play a measurable role in reading out scene identity.

We additionally investigate the functional effects of the variation in couplings across movies. To do this, we calculate how quickly evidence accumulates for the correct scene in the counterfactual case where the couplings between movies have been swapped between scenes. In that chimera case, we again find that the performance is identical to that of the full model (Fig. 4B, dotted red line).

**Couplings Arise from Both Gap Junctions and Shared Bipolar Cell Input.** We have found consistent population structure that is dominated by sparse couplings. This structure is hard-wired into the retina code and could arise from many different local



**Fig. 4.** The conserved, sparse coupling structure contributes to readout of scene identity. (A) Cartoon of how we consider the decoding problem downstream of the retina. Signals pass to the retina and are encoded, they then must be decoded downstream. Here, we will consider the task of decoding scene identity as a proxy for decoding large-scale scene statistics. (B) Decoding performance of a variety of models, as measured in number of seconds to achieve high certainty about the correct scene identity. There is a significant difference between the performance of networks with or without couplings, but no difference between networks with the full set of couplings, with only the skeleton couplings, and networks with swapped couplings.

circuit motifs. Pairwise retinal couplings could be correlated with shared upstream input from a bipolar cell (Fig. 5A, Left), direct gap junctions between RGCs (Fig. 5A, Top), or gap junctions between RGCs and a third neuron such as an amacrine cell (Fig. 5A, Right). All of these mechanisms have been previously implicated in phenomena that could play a significant role in computation in natural scenes. Lateral connectivity (through amacrine cells or gap junctions) is involved in many anticipatory behaviors in the retina (65–67), while feedforward input from bipolar cells to RGCs is involved in many adaptation phenomena (68, 69).

Nonlinear summation of bipolar cell (BC) inputs has been shown to be an integral component of retinal computation (70–75). The convergence of BCs onto RGCs is modeled using nonlinear summation in so-called cascade models. These models explain a wide array of complex retinal computations [e.g., motion onset (70), omitted stimulus response (76), background vs. object motion (77), reversal response (78–80)]. As a complement to that, BCs have diverging projections onto multiple RGCs on the retina (81). While bipolar cells sample a smaller portion of the visual scene than RGCs, gap junction networks between RGCs can expand the bipolar projective field out as far as  $\sim 1$  mm (82). This BC divergence could play a role in shaping the population code in the retina and needs to be further explored with naturalistic, dynamic inputs.

In order to detect putative bipolar cell inputs onto the RGCs in our dataset, we use spike-triggered independent component analysis (ST-ICA) (83) on the white noise checkerboard stimulus. ST-ICA models each RGC as the output of a temporal filter and spatial subunits. Similarly, these subunits have been experimentally shown to map to bipolar cell inputs (84). We

show an example of the spatial subunits and temporal filters for three fast-OFF RGCs in Fig. 5B and C, two pairs of which (blue, pink; green, pink) demonstrate strong couplings. One coupling pair (blue, pink) exhibits a highly overlapped spatial subunit that we classify as a shared upstream input (*Materials and Methods*). In Fig. 5C, temporal filters for two cells (blue, green) show nearly identical characteristics. This may arise from a gap junction connecting them and might explain their strong coupling (50).

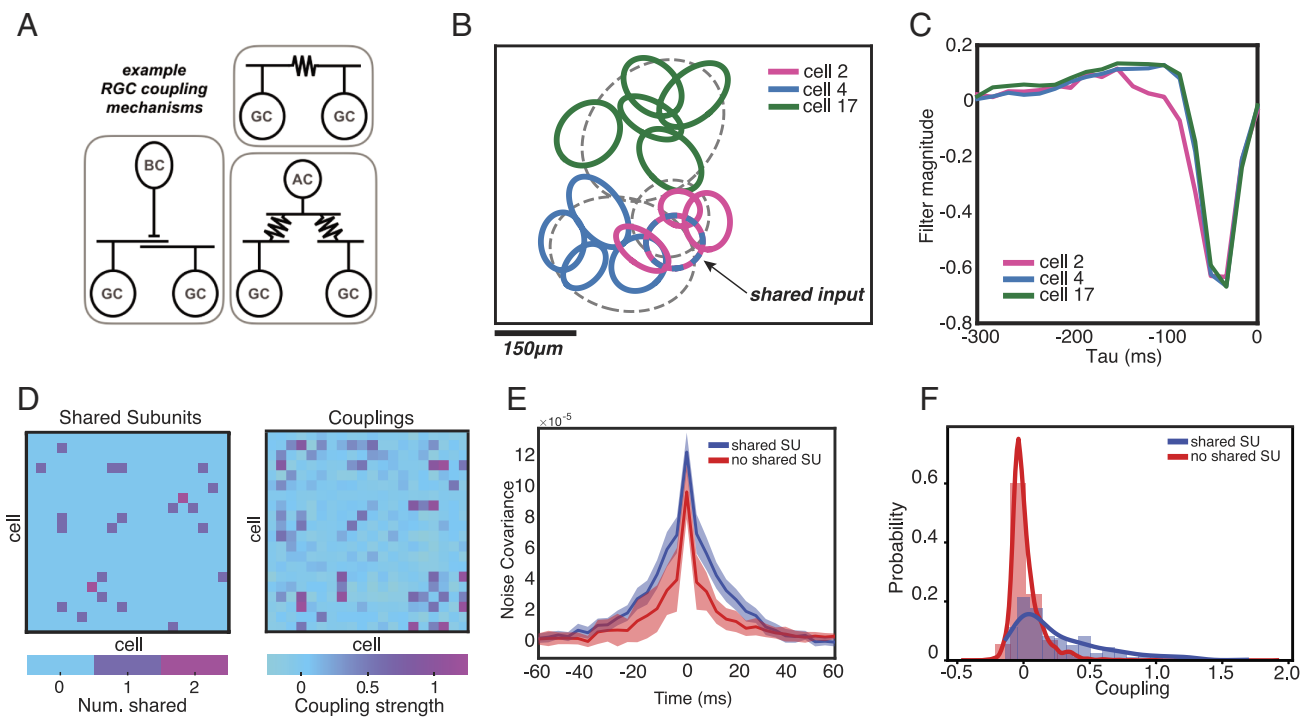
Across the population, many RGCs share spatial subunits. These subunits closely align with the observed coupling matrix from the stimulus-dependent maximum entropy models (Fig. 5D), demonstrating that strong couplings may arise in part from shared upstream input. However, not all cells with strong couplings share upstream input, and the presence of a shared subunit alone does not guarantee the existence of a strong coupling. Thus, strong couplings might arise from multiple sources within the retinal cell population.

Previous work has suggested that gap junctions may underpin couplings between RGCs (8, 50). Here, we find evidence of gap junctions by inspecting the cross-covariance of responses after subtracting the trial-average (at zero lag this is the usual noise correlation). Pairwise noise correlations due to gap junctions can generally be split into two classes, direct RGC–RGC couplings and shared gap junctions with a third upstream neuron. The symmetric, medium width correlation we observe between some highly coupled cells without a shared subunit (Fig. 5D, red) likely arises from this second class, as direct RGC–RGC gap junctions lead to transient noise correlations at a submillisecond timescale (8). Furthermore, the broader noise correlations between RGCs with shared subunits demonstrate a timescale longer than can be explained from gap junctions alone and may indicate a coupling arising from shared upstream input (Fig. 5A, blue).

Shared input between RGCs, whatever the source, greatly increases the likelihood of a strong coupling compared to RGC pairs without shared input (Fig. 5E). Of course, shared input does not *guarantee* a strong coupling between RGCs, but the long tail of strong couplings for pairs with shared input suggests that these mechanisms underpin our sparse network of strong interactions. We hypothesize that shared bipolar input and gap junctions work in concert to generate a sparse set of intrinsic correlations between RGCs.

## Discussion

Sparse, strong couplings between cells in this neural population are an important component of both generating the population response repertoire and potential downstream readout of natural scene identity. While some studies show that independent models of the retinal code retain upward of 90% of the response structure (85), such analyses are typically agnostic to the downstream readout goals of the organism. Without a defined readout goal, it is impossible to determine whether aspects of the response structure lost by an independent encoding relay information meaningful to the organism. It is possible that an independent readout preserves much of the population response range while failing to effectively convey critical features of the visual scene. Natural scenes probe a behaviorally relevant context to assess the impact of noise correlations on neural coding. These movies, like other natural inputs, drive a richer and more reliable code in the brain (86–88). Comparing across movies reveals what the more subtle features in the neural code might be used for.



**Fig. 5.** Strong couplings result from both shared inputs and gap junctions. (A) Three potential pairwise coupling sources. *Left*, shared input from an upstream bipolar cell. *Top*, a direct gap junction between two RGCs. *Bottom Right*, gap junction connections to a shared third neuron, an amacrine cell. (B) ST-ICA modeled spatial subunits for three cells, two of whom share a subunit. (C) Time filters for the same cells shown in B. (D) A comparison between shared subunits and couplings indicates many highly coupled cells shared upstream input. (E) The average noise covariance for highly coupled cells. Highly coupled cells that share a subunit have a broader noise covariance in time than coupled cells without a shared subunit. (F) Distribution of coupling strengths, both for cells sharing subunits and those that do not.

Both cell–cell couplings and time-dependent fluctuations in individual cell excitability shape the observed retinal response distributions and convey information about scene features. Previous work has shown that an independent model of retinal activity can recover information about scene features, like spatial texture and local motion (89). This study used the trial averaged, time-dependent firing rates across the full retinal population to show how a lower-dimensional, latent representation of this activity could support precise readout of time in a natural movie. These latent representations also generalized across different natural movies [(89), in figure 3A]. The local time-dependent fields in our maximum entropy model are the key features for this kind of scene encoding, while our sparse skeleton of interactions encodes broader categorical scene information. The interplay between these static and dynamic features provides a rich substrate for encoding different kinds of scene information flexibly and reliably.

Only a sparse subset of strong couplings is critical for scene identity decoding. There is a distinction between this kind of “functional sparsity” (where the couplings matter for performance) and “statistical sparsity” (where the distribution is, itself, sparse). Previous work has shown that cortical populations are both statistically and functionally sparse (90). This could be a general strategy by which different kinds of information are multiplexed in a neural code (91, 92).

Our finding that sparse interactions sufficiently capture the functional impact of noise correlations on neural encoding elaborates on previous work that argued for a dense network of weak couplings in the retinal code (37, 41). In these early models, both the fields and interactions between cells were static. The fluctuating fields we included (following the time-dependent

maximum entropy model proposed in ref. 50) absorb much of that structure, and capture the independent component of stimulus-driven changes in the neural population response. The remaining sparse couplings are potentially the key factor for efficient scene identification. A sparse skeleton of interactions may be flexibly activated in different scenes in a way that is easier to discriminate downstream. These interactions generate noise correlations that may reflect changes in external scene correlation structure, which may help recover scene-specific information that is otherwise lost to single-cell-level adaptation (26). Our finding that notably, these interactions allow for scene decoding at a similar timescale as some behavioral features in amphibians (93–95).

On the flip side, sparse codes might hamstring error correction (96, 97), so future work should explore how these costs and benefits trade-off for behaviorally relevant inputs and tasks. New holography work may provide the kind of detailed probe needed to explore the dual roles of individual cell responses and cell–cell correlations in natural tasks (98–100).

These data provide some hints at the mechanisms for generating this sparse but strong functional connectivity in the retina. In many ways, the circuit structure in the eye differs from that found in the cortex, yet some common anatomic motifs are found in both. The retina is not a recurrent neural network, RGCs do not have direct synaptic coupling, and the photoreceptor-to-RGC circuit is largely feed-forward. However, gap junction coupling is prevalent throughout the brain, as is shared, convergent input to regional projection neurons. To create a population code with sparse interactions, the retina needs to be wired around its particular structural constraints, but appears to use these common motifs found throughout the brain. Our initial results

indicate that sparse interactions may be the result of common bipolar inputs and gap junction coupling between RGCs. Both gap junctions and common bipolar inputs lead to stronger coupling between cells, but our analysis is not sensitive enough to tease apart whether these two types of coupling sources are mutually exclusive in particular pairs of cells. Exclusivity would be an efficient way to implement a sparse skeleton of specific cell–cell interactions. Future work to disentangle the circuit mechanisms giving rise to the sparse skeleton might ultimately inform studies in the cortex where gap junction coupling is also present (101–103).

## Materials and Methods

**Neural Data.** Voltage traces from the output, RGC layer of a larval tiger salamander retina were recorded following the methods outlined in ref. 17. In brief, the retina was isolated in darkness and pressed against a 252-channel multielectrode array. Voltage recordings were taken during stimulus presentation of both natural movies and white noise stimuli, and spike-sorted using an automated clustering algorithm that was hand curated after initial template clustering and fits. This technique captured a highly overlapping neural population of 93 cells that fully tiled the recorded region of visual space. Spike times were binned at 16.667 ms for all analyses presented.

The data has been published at DOI: [10.5061/dryad.4qrfj6qmq8](https://doi.org/10.5061/dryad.4qrfj6qmq8) (104).

**Visual Stimuli.** A white noise checkerboard stimulus (with binary white and black squares) was played at 30 frames per second (fps) for 30 min prior to and after the natural scene stimuli. Five different natural movies lasting 20 s were played in a pseudorandom order, and each was displayed a minimum of 80 times. The movies labeled tree, water, grasses, fish, and self-motion were repeated 83, 80, 84, 91, and 85 times, respectively. All natural scenes except for the tree stimulus were displayed at 60 fps. The tree stimulus was updated at a rate of 30 fps with each frame repeated twice to match the 60 fps frame rate of the other movies.

In all movies, the cells significantly increase their firing rates in the first 200 ms (Fig. 2A) following the switch to a new stimulus. This is followed by a rapid decay back to a baseline firing rate. This is likely due to a strong population response to abrupt changes in luminance within their receptive fields (68, 105, 106). In subsequent analysis, we exclude the first 500 ms of every trial to isolate the more steady-state response of the retina to scene-specific features and dynamics.

**Maximum Entropy Modeling.** We followed the data-driven algorithm introduced in ref. 107 for our maximum entropy modeling. This algorithm uses a quasi-Newtonian method that allows for inference of model parameters  $\mathbf{X}$ , in our case, time-dependent fields  $h_i^t$  and constant couplings  $J_{ij}$ , without needing to compute the inverse model susceptibility matrix  $\chi^{-1}[\mathbf{X}]$  at each time point during the learning dynamics. As in ref. 50, we learn a maximum entropy model with time-varying fields. Specifically, we learn a model of form

$$P(\sigma_1^t \dots \sigma_N^t) \propto \exp \left( - \sum_i^N h_i^t \sigma_i^t - \sum_{i < j}^N J_{ij} \sigma_i^t \sigma_j^t \right). \quad [3]$$

The time-dependent fields  $h_i^t$  capture the time-varying firing rate of cells  $\sigma_i$  in response to the stimulus. This means that stimulus-dependent correlations between cells are absorbed into the fields, leaving the couplings  $J_{ij}$  to encapsulate the noise correlations between cells.

All fits were done on groups of 20 cells, which were all subsets of the full population of 93 cells. These subsets were chosen at random.

We additionally used data from sparse models to generate the skeleton and independent information shown in Fig. 3. In order to train the independent model, we follow the same procedure as our full pairwise models, but we constrain all couplings to be zero at all fitting steps. In order to train the skeleton model, we first fit a full model with all couplings and fields. We then select

the top 10% of couplings, and refit a model where the only allowed nonzero couplings are pairs identified in those top 10% of couplings.

Model performance was tested by comparing predicted probabilities of the population response states to observed frequencies in the data. This is a dramatic increase in the number of parameters from  $N^2$  to  $2^N$ , and represents a difficult extrapolation task.

**DCA Models.** When we were looking for a model that would give us access to the couplings in data from the white noise stimulus, we were faced with the lack of stimulus repeats. The stimulus-dependent maximum entropy model fundamentally relies on measurements of noise correlations, and as such does not work without stimulus repeats. However, we also wished for a method that would allow us to infer couplings that are only representative of the noise correlations, not couplings that included both shared stimulus inputs and noise correlations. In general, this is not possible, and standard maximum entropy models are not up to the task. Here, however, we took advantage of a feature of these data: many of our cells have highly overlapping receptive fields. This means that two cells with similar spatial receptive fields can be correlated two ways: because of a noise correlation, or because of shared stimulus drive. In particular, if several cells are driven in the same way by the stimulus, they will all have correlations with each other, and therefore many methods would infer “loops” of couplings.

However, in the protein community, DCA was developed as a maximum-entropy technique with an emphasis on ignoring indirect couplings (i.e., breaking loops) (54, 108), and a prior on the sparsity of the coupling matrix. This means that in our case, with several cells all driven locally in the same way by the stimulus, many of those correlations will be dropped in favor of a sparser explanation for the population activity. In particular, one should expect the remaining couplings to be the strongest correlations—those where there is both a biological coupling between the cells and a shared receptive field. This means that we do not expect quantitative agreement between this method and stimulus-dependent maximum entropy but that we can hope to find the same skeleton of strong couplings.

In order to fit the DCA model, we followed methods discussed in ref. 54. We chose a gauge where neural silence or activity are described by  $\{0,1\}$  to more easily relate the DCA network to couplings fit from the time-dependent maxent models.

**Decoding Scene Identity.** For the scene identity decoding introduced in Fig. 3B, we used a Bayesian approach to measure  $P(\text{movie}|\text{spikes})$ ,

$$P(\text{movie}|\text{spikes}) = \frac{P(\text{spikes}|\text{movie})P(\text{movie})}{P(\text{spikes})}. \quad [4]$$

This describes the decoding task—deciding which of the five movies is being shown, based on the population activity from 20 neurons. This is highly related to the problem solved by downstream networks reading out retinal output and attempting to determine large-scale information about the environment. In this case, the decision is between five equi-probable scenes, so we have a uniform prior,  $P(\text{movie}) = 1/5$ .

In order to calculate  $P(\text{spikes}|\text{movie})$ , we need to sample spike trains from our models. Our goal is to evaluate the ability of various underlying generative architectures to support scene discriminability. Sampling from the models yields two key advantages: first, we can generate as many samples as needed from those distributions, and second, we can directly investigate the effects of degrading couplings or swapping couplings between scenes. “Spike data” for the model will be a set of vectors  $\{\vec{\sigma}_i\}$ , where each is the binary representation of spiking of all neurons in time bin  $i$ . The  $j$ 'th element of  $\vec{\sigma}_i$  represents neuron  $j$ 's activity in time bin  $i$ . All of our models produce a time-dependent probability distribution of joint activity of the neurons,  $P(\vec{\sigma}|t)$ , but also produces a time-independent distribution,  $P(\vec{\sigma}) = \sum_t (1/T)P(\vec{\sigma}|t)$ . We will generate each spike vector  $\vec{\sigma}$  by pulling from this time-independent distribution. We do this separately for each time bin in our simulated spike train.

Armed with a set of neural responses to a particular movie  $\{\vec{\sigma}_i\}$ , we then can calculate  $P(\vec{\sigma}_i|\text{movie})$ . For movie  $\alpha$ , we calculate  $P(\vec{\sigma}_i|\text{movie}_\alpha)$  based on each model.

To calculate  $P(\{\bar{\sigma}_j\}|\text{movie}_\alpha)$  at time point  $\tau$  on the plot in Fig. 4B, we need to accumulate evidence,  $P(\{\bar{\sigma}_1, \bar{\sigma}_2, \dots, \bar{\sigma}_\tau\}|\text{movie}_\alpha) = \prod_{j=1}^{\tau} P(\bar{\sigma}_j|\text{movie}_\alpha)$ . Finally, we can calculate the overall probability of a given spike pattern,  $P(\{\bar{\sigma}_j\})$ , by summing over movies,  $P(\{\bar{\sigma}_j\}) = \sum_{\alpha} P(\bar{\sigma}_j|\text{movie}_\alpha)$ . We can easily change the underlying models and resample from the population patterns, allowing for easy comparisons on this decoding task.

In order to generate Fig. 4B, for each combination of movie and model choice, we generate 1,000 spike trains. We show the median performance of each model at decoding, and we show the quartiles for the full and conditionally independent models.

**Spike-Triggered ICA.** To compute the Spike-Triggered ICA (ST-ICA) we follow methods developed in ref. 83.

We first compute the spike-triggered average for each cell in a natural cubic spline basis. This is a common method to reduce the number of parameters needed for the model and ensure that the resulting receptive fields are smooth in space and time. We choose the number of splines such that the log-likelihood on held-out white noise data is maximized.

To further reduce the number of parameters, we assume our receptive field is rank 1, it can be separated into a spatial filter and a temporal filter. Following these assumptions, we use SVD to find this rank 1 approximation. This provably minimizes reconstruction error under the Frobenius norm. We then crop the spatial dimensions for each cell to the regions containing the receptive fields and convolve the stimulus with the temporal filters, which leaves only spatial degrees of freedom.

For each cell, we then have a matrix of size  $N$  spikes by  $M$  features, where each feature is a spatial pixel convolved with time filter. We use Preconditioned ICA

(109), an algorithm for ICA that uses preconditioned L-BFGS, a low-memory quasi-Newton optimization algorithm, for optimization to estimate 20 independent components. Resulting components were considered proper subunit candidates based on the presence of significant spatial autocorrelations, following methods in ref. 84.

With a list of candidate subunits for each cell, we then compute the activation of that subunit by projecting the time-convolved stimulus onto each filter identified by ST-ICA. Two units are considered the same following methods developed in ref. 110.

**Data, Materials, and Software Availability.** The data has been published at <https://doi.org/10.5061/dryad.4qrj6qm8> (104).

**ACKNOWLEDGMENTS.** We thank Ulisse Ferrari for helpful discussions and providing a preliminary version of the code used for the time-dependent maximum entropy models. This work was supported in part by the NSF, through the Center for the Physics of Biological Function (PHY-1734030) and an NSF Graduate Research Fellowship (DGE-2039656) (CMH); and by the NIH BRAIN initiative (R01EB026943).

Author affiliations: <sup>a</sup>Committee on Computational Neuroscience, Department of Organismal Biology and Anatomy, University of Chicago, Chicago, IL 60637; <sup>b</sup>Department of Physics, Princeton University, Princeton, NJ 08540; <sup>c</sup>Department of Organismal Biology and Anatomy, University of Chicago, Chicago, IL 60637; <sup>d</sup>Department of Physics, University of Chicago, Chicago, IL 60637; <sup>e</sup>Princeton Neuroscience Institute, Department of Molecular Biology, Princeton University, Princeton, NJ 08540; <sup>f</sup>Institut de la Vision, Sorbonne Université, INSERM, Paris 75012, France; and <sup>g</sup>Center for the Physics of Biological Function, Department of Physics, Princeton University, Princeton, NJ 08540

- H. B. Barlow, Summation and inhibition in the frog's retina. *J. Physiol.* **119**, 69 (1953).
- D. H. Hubel, T. N. Wiesel, Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* **148**, 574 (1959).
- H. K. Hartline, The receptive fields of optic nerve fibers. *Am. J. Physiol.-Legacy Content* **130**, 690-699 (1940).
- D. K. Warland, P. Reinagel, M. Meister, Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* **78**, 2336-2350 (1997).
- W. Bialek, F. Rieke, R. van Steveninck, D. Warland, Reading a neural code. *Adv. Neural Inf. Process. Syst.* **2**, 36-43 (1989).
- T. Gollisch, M. Meister, Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron* **65**, 150-164 (2010).
- S. A. Baccus, B. P. Ölveczky, M. Manu, M. Meister, A retinal circuit that computes object motion. *J. Neurosci.* **28**, 6807-6817 (2008).
- I. H. Brivanlou, D. K. Warland, M. Meister, Mechanisms of concerted firing among retinal ganglion cells. *Neuron* **20**, 527-539 (1998).
- J. Shlens, F. Rieke, E. Chichilnisky, Synchronized firing in the retina. *Curr. Opin. Neurobiol.* **18**, 396-402 (2008).
- J. W. Pillow *et al.*, Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995-999 (2008).
- L. Ramirez, W. Bialek, Compression as a path to simplification: Models of collective neural activity. *arXiv [Preprint]* (2021). <https://arxiv.org/abs/2112.14334> (Accessed 29 November 2024).
- D. B. Kastner, S. A. Baccus, T. O. Sharpee, Critical and maximally informative encoding between neural populations in the retina. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 2533-2538 (2015).
- N. Maheswaranathan *et al.*, Interpreting the retinal neural code for natural scenes: From computations to neurons. *Neuron* **111**, 2742-2755 (2023).
- V. Botella-Soler, S. Deny, G. Martius, O. Marre, G. Tkačik, Nonlinear decoding of a complex movie from the mammalian retina. *PLoS Comput. Biol.* **14**, e1006057 (2018).
- M. Molano-Mazon, A. Onken, E. Piasini, S. Panzeri, "Synthesizing realistic neural population activity patterns using Generative Adversarial Networks" in *International Conference on Learning Representations* (2018). <https://openreview.net/forum?id=r1VvseBzAZ>. Accessed 29 November 2024.
- C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini, K. D. Harris, High-dimensional geometry of population responses in visual cortex. *Nature* **571**, 361-365 (2019).
- O. Marre *et al.*, Mapping a complete neural population in the retina. *J. Neurosci.* **32**, 14859-14873 (2012).
- A. Berényi *et al.*, Large-scale, high-density (up to 512 channels) recording of local circuits in behaving animals. *J. Neurophysiol.* **111**, 1132-1149 (2014).
- C. M. Lopez *et al.*, "22.7 a 966-electrode neural probe with 384 configurable channels in 0.13 mm SOI CMOS" in *2016 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, 2016), pp. 392-393.
- N. A. Steinmetz *et al.*, Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* **372**, eabf4588 (2021).
- R. W. Rodieck, *The First Steps in Seeing* (Sinauer Associates, 1998).
- D. L. Ruderman, W. Bialek, Statistics of natural images: Scaling in the woods. *Phys. Rev. Lett.* **73**, 814 (1994).
- O. Schwartz, E. P. Simoncelli, Natural signal statistics and sensory gain control. *Nat. Neurosci.* **4**, 819-825 (2001).
- D. W. Dong, J. J. Atick, Statistics of natural time-varying images. *Netw. Comput. Neural Syst.* **6**, 345 (1995).
- J. H. van Hateren, D. L. Ruderman, Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc. R. Soc. Lond. B* **265**, 2315-2320 (1998).
- A. L. Fairhall, G. D. Lewen, W. Bialek, R. R. de Ruyter van Steveninck, Efficiency and ambiguity in an adaptive neural code. *Nature* **412**, 787-792 (2001).
- J. J. Atick, A. N. Redlich, What does the retina know about natural scenes? *Neural Comput.* **4**, 196-210 (1992).
- I. Nemenman, G. D. Lewen, W. Bialek, R. R. de Ruyter van Steveninck, Neural coding of natural stimuli: Information at sub-millisecond resolution. *PLoS Comput. Biol.* **4**, e1000025 (2008).
- J. Jovancevic-Misic, M. Hayhoe, Adaptive gaze control in natural environments. *J. Neurosci.* **29**, 6234-6238 (2009).
- M. J. Zimmermann *et al.*, Zebrafish differentially process color across visual space to match natural scenes. *Curr. Biol.* **28**, 2018-2032.e5 (2018).
- M. Wong-Riley, Synaptic organization of the inner plexiform layer in the retina of the tiger salamander. *J. Neurocytol.* **3**, 1-33 (1974).
- D. A. Burkhardt, P. K. Fahey, M. A. Sikora, Natural images and contrast encoding in bipolar cells in the retina of the land-and aquatic-phase tiger salamander. *Vis. Neurosci.* **23**, 35-47 (2006).
- J. M. Salisbury, S. E. Palmer, A dynamic scale-mixture model of motion in natural scenes. *bioRxiv [Preprint]* (2023). <https://doi.org/10.1101/2023.10.19.563101> (Accessed 29 November 2024).
- J. M. Salisbury, S. E. Palmer, Optimal prediction in the retina and natural motion statistics. *J. Stat. Phys.* **162**, 1309-1323 (2016).
- K. J. Kim, F. Rieke, Slow Na<sup>+</sup> inactivation and variance adaptation in salamander retinal ganglion cells. *J. Neurosci.* **23**, 1506-1516 (2003).
- F. Rieke, Temporal contrast adaptation in salamander bipolar cells. *J. Neurosci.* **21**, 9445-9454 (2001).
- E. Schneidman, M. J. Berry, R. Segev, W. Bialek, Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007-1012 (2006).
- J. W. Pillow, E. P. Simoncelli, Dimensionality reduction in neural models: An information-theoretic generalization of spike-triggered average and covariance analysis. *J. Vis.* **6**, 9 (2006).
- E. Granot-Atedgi, G. Tkačik, R. Segev, E. Schneidman, Stimulus-dependent maximum entropy models of neural population codes. *PLoS Comput. Biol.* **9**, e1002922 (2013).
- E. Ganmor, R. Segev, E. Schneidman, Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 9679-9684 (2011).
- G. Tkačik *et al.*, Searching for collective behavior in a large network of sensory neurons. *PLoS Comput. Biol.* **10**, e1003408 (2014).
- Y. Roudi, S. Nirenberg, P. E. Latham, Pairwise maximum entropy models for studying large biological systems: When they can work and when they can't. *PLoS Comput. Biol.* **5**, e1000380 (2009).
- E. T. Jaynes, Information theory and statistical mechanics. *Phys. Rev.* **106**, 620 (1957).
- G. Tkačik, J. S. Prentice, V. Balasubramanian, E. Schneidman, Optimal population coding by noisy spiking neurons. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 14419-14424 (2010).
- S. W. Linderman, R. P. Adams, J. W. Pillow, Inferring structured connectivity from spike trains under negative-binomial generalized linear models. *Neuron* **50**, 150 (2015).

46. S. L. Keeley, D. M. Zoltowski, M. C. Aoi, J. W. Pillow, Modeling statistical dependencies in multi-region spike train data. *Curr. Opin. Neurobiol.* **65**, 194–202 (2020).
47. S. Gerwinn, J. H. Macke, M. Bethge, Bayesian inference for generalized linear models for spiking neurons. *Front. Comput. Neurosci.* **4**, 1299 (2010).
48. F. Gerhard, M. Deger, W. Truccolo, On the stability and dynamics of stochastic spiking neuron models: Nonlinear Hawkes process and point process GLMs. *PLoS Comput. Biol.* **13**, e1005390 (2017).
49. J. Shlens *et al.*, The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.* **26**, 8254–8266 (2006).
50. U. Ferrari *et al.*, Separating intrinsic interactions from extrinsic correlations in a network of sensory neurons. *Phys. Rev. E* **98**, 042410 (2018).
51. L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, W. Bialek, Successes and failures of simplified models for a network of real neurons. arXiv [Preprint] (2021). <https://arxiv.org/abs/2112.14735> (Accessed 29 November 2024).
52. O. Sorochnytskiy, S. Deny, O. Marre, U. Ferrari, Predicting synchronous firing of large neural populations from sequential recordings. *PLoS Comput. Biol.* **17**, e1008501 (2021).
53. K. D. Simmons *et al.*, Transformation of stimulus correlations by the retina. *PLoS Comput. Biol.* **9**, e1003344 (2013).
54. M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, T. Hwa, Identification of direct residue contacts in protein-protein interaction by message passing. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 67–72 (2009).
55. G. Buzsáki, K. Mizuseki, The log-dynamic brain: How skewed distributions affect network operations. *Nat. Rev. Neurosci.* **15**, 264–278 (2014).
56. S. Ganguli, H. Sompolinsky, Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annu. Rev. Neurosci.* **35**, 485–508 (2012).
57. S. Ebrahimi *et al.*, Emergent reliability in sensory cortical coding and inter-area communication. *Nature* **605**, 713–721 (2022).
58. E. J. Candès, J. K. Romberg, T. Tao, Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.* **59**, 1207–1223 (2006).
59. T. Watanabe *et al.*, A pairwise maximum entropy model accurately describes resting-state human brain networks. *Nat. Commun.* **4**, 1370 (2013).
60. I. E. Ohiorhenuan *et al.*, Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* **466**, 617–621 (2010).
61. S. I. Firth, C. T. Wang, M. B. Feller, Retinal waves: Mechanisms and function in visual system development. *Cell Calcium* **37**, 425–432 (2005).
62. J. Cafaro, F. Rieke, Noise correlations improve response fidelity and stimulus encoding. *Nature* **468**, 964–967 (2010).
63. D. Lee, J. Kim, S. A. Baccus, Classification and analysis of retinal interneurons by computational structure under natural scenes. bioRxiv [Preprint] (2024). <https://doi.org/10.1101/2024.03.18.585364> (Accessed 29 November 2024).
64. B. B. Averbeck, P. E. Latham, A. Pouget, Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
65. S. Souhail, H. Cessac, On the potential role of lateral connectivity in retinal anticipation. *J. Math. Neurosci.* **11**, 1–60 (2021).
66. S. Trenholm, A. J. McLaughlin, D. J. Schwab, G. B. Awatramani, Dynamic tuning of electrical and chemical synaptic transmission in a network of motion coding retinal neurons. *J. Neurosci.* **33**, 14927–14938 (2013).
67. S. Trenholm, D. J. Schwab, V. Balasubramanian, G. B. Awatramani, Lag normalization in an electrically coupled neural network. *Nat. Neurosci.* **16**, 154–156 (2013).
68. T. Jarsky *et al.*, A synaptic mechanism for retinal adaptation to luminance and contrast. *J. Neurosci.* **31**, 11003–11015 (2011).
69. M. B. Manookin, J. B. Demb, Presynaptic mechanism for slow contrast adaptation in mammalian retinal ganglion cells. *Neuron* **50**, 453–464 (2006).
70. E. Y. Chen *et al.*, Alert response to motion onset in the retina. *J. Neurosci.* **33**, 120–132 (2013).
71. M. A. Goldin *et al.*, Context-dependent selectivity to natural images in the retina. *Nat. Commun.* **13**, 5556 (2022).
72. T. Gollisch, Features and functions of nonlinear spatial integration by retinal ganglion cells. *J. Physiol.-Paris* **107**, 338–348 (2013).
73. G. Schwartz, F. Rieke, Nonlinear spatial encoding by retinal ganglion cells: When  $1 + 1 \neq 2$ . *J. Gen. Physiol.* **138**, 283–290 (2011).
74. A. L. Fairhall *et al.*, Selectivity for multiple stimulus features in retinal ganglion cells. *J. Neurophysiol.* **96**, 2724–2738 (2006).
75. S. J. Zapp, S. Nitsche, T. Gollisch, Retinal receptive-field substructure: Scaffolding for coding and computation. *Trends Neurosci.* **45**, 430–445 (2022).
76. G. Schwartz, R. Harris, D. Shrom, M. J. Berry, Detection and prediction of periodic patterns by the retina. *Nat. Neurosci.* **10**, 552–554 (2007).
77. B. P. Ölveczky, S. A. Baccus, M. Meister, Segregation of object and background motion in the retina. *Nature* **423**, 401–408 (2003).
78. G. Schwartz, S. Taylor, C. Fisher, R. Harris, M. J. Berry, Synchronized firing among retinal ganglion cells signals motion reversal. *Neuron* **55**, 958–969 (2007).
79. N. P. Shah *et al.*, Inference of nonlinear receptive field subunits with spike-triggered clustering. *eLife* **9**, e45743 (2020).
80. E. Y. Chen, J. Chou, J. Park, G. Schwartz, M. J. Berry, The neural circuit mechanisms underlying the retinal response to motion reversal. *J. Neurosci.* **34**, 15557–15575 (2014).
81. H. Asari, M. Meister, Divergence of visual channels in the inner retina. *Nat. Neurosci.* **15**, 1581–1589 (2012).
82. H. Asari, M. Meister, The projective field of retinal bipolar cells and its modulation by visual context. *Neuron* **81**, 641–652 (2014).
83. A. Saleem, H. Krapp, S. R. Schultz, Receptive field characterization by spike-triggered independent component analysis. *J. Vis.* **8**, 2.1–2.16 (2008).
84. J. K. Liu *et al.*, Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nat. Commun.* **8**, 1–14 (2017).
85. S. Nirenberg, S. M. Carcieri, A. L. Jacobs, P. E. Latham, Retinal ganglion cells act largely as independent encoders. *Nature* **411**, 698–701 (2001).
86. R. V. Rikhye, M. Sur, Spatial correlations in natural scenes modulate response reliability in mouse visual cortex. *J. Neurosci.* **35**, 14661–14680 (2015).
87. E. Froudarakis *et al.*, Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat. Neurosci.* **17**, 851–857 (2014).
88. U. Hasson, R. Malach, D. J. Heeger, Reliability of cortical activity during natural stimulation. *Trends Cogn. Sci.* **14**, 40–48 (2010).
89. S. Wang *et al.*, Learning low-dimensional generalizable natural features from retina using a U-net. *Adv. Neural Inf. Process. Syst.* **35**, 11355–11368 (2022).
90. K. Bojanek, Y. Zhu, J. MacLean, Cyclic transitions between higher order motifs underlie sustained asynchronous spiking in sparse recurrent networks. *PLoS Comput. Biol.* **16**, e1007409 (2020).
91. S. Panzeri, N. Brunel, N. K. Logothetis, C. Kayser, Sensory neural codes using multiplexed temporal scales. *Trends Neurosci.* **33**, 111–120 (2010).
92. R. Naud, H. Sprekeler, Sparse bursts optimize information transmission in a multiplexed neural code. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E6329–E6338 (2018).
93. K. Brändle, Quantitative studies of the reactions to horizontal angular accelerations in axolotls: I. The head-turning reflexes of normal animals. *J. Exp. Biol.* **66**, 1–14 (1977).
94. K. Brändle, Quantitative studies of the reactions to horizontal angular accelerations in axolotls: II. Head-turning reflexes in animals with a supernumerary pair of labyrinths. *J. Exp. Biol.* **66**, 15–31 (1977).
95. J. Kopp, G. Manteuffel, Quantitative analysis of salamander horizontal head nystagmus. *Brain Behav. Evol.* **25**, 187–196 (1984).
96. J. L. Puchalla, E. Schneidman, R. A. Harris, M. J. Berry, Redundancy in the population code of the retina. *Neuron* **46**, 493–504 (2005).
97. E. Gannor, R. Segev, E. Schneidman, A thesaurus for a neural population code. *eLife* **4**, e06134 (2015).
98. H. Adesnik, L. Abdeladim, Probing neural codes with two-photon holographic optogenetics. *Nat. Neurosci.* **24**, 1356–1366 (2021).
99. M. Dal Maschio, J. C. Donovan, T. O. Helmbrecht, H. Baier, Linking neurons to network function and behavior by two-photon holographic optogenetics and volumetric imaging. *Neuron* **94**, 774–789 (2017).
100. W. Yang, L. Carrillo-Reid, Y. Bando, D. S. Peterka, R. Yuste, Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions. *eLife* **7**, e32671 (2018).
101. D. S. Friend, N. B. Gilula, Variations in tight and gap junctions in mammalian tissues. *J. Cell Biol.* **53**, 758–776 (1972).
102. A. Peinado, R. Yuste, L. C. Katz, Gap junctional communication and the development of local circuits in neocortex. *Cereb. Cortex* **3**, 488–498 (1993).
103. Y. Li *et al.*, Clonally related visual cortical neurons show similar stimulus feature selectivity. *Nature* **486**, 118–121 (2012).
104. S. E. Palmer *et al.*, Larval salamander retinal population data in response to natural movies from the Chicago Motion Database [Dataset]. Dryad. <https://doi.org/10.5061/dryad.4qrfj6qm8>. Deposited 29 November 2024.
105. J. B. Demb, Functional circuitry of visual adaptation in the retina. *J. Physiol.* **586**, 4377–4384 (2008).
106. K. Nakatani, K. W. Yau, Calcium and light adaptation in retinal rods and cones. *Nature* **334**, 69–71 (1988).
107. U. Ferrari, Learning maximum entropy models from finite-size data sets: A fast data-driven algorithm allows sampling from the posterior distribution. *Phys. Rev. E* **94**, 023301 (2016).
108. F. Morcos *et al.*, Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proc. Natl. Acad. Sci. U.S.A.* **108**, E1293–E1301 (2011).
109. P. Ablin, J. F. Cardoso, A. Gramfort, Faster independent component analysis by preconditioning with Hessian approximations. *IEEE Trans. Signal Process.* **66**, 4040–4049 (2018).
110. S. Jia *et al.*, Neural system identification with spike-triggered non-negative matrix factorization. *IEEE Trans. Cybern.* **52**, 4772–4783 (2021).