doi: 10.1038/nature07140

Spatio-temporal correlations and visual signaling in a complete neuronal population

Jonathan W. Pillow[1], Jonathon Shlens[2], Liam Paninski[3],
Alexander Sher[4], Alan M. Litke[4], E. J. Chichilnisky[2], Eero P. Simoncelli[5]

[1] Gatsby Computational Neuroscience Unit, UCL
[2] The Salk Institute, La Jolla, California
[3] Department of Statistics and Center for Theoretical Neuroscience, Columbia University
[4] Santa Cruz Institute for Particle Physics, University of California, Santa Cruz
[5] Howard Hughes Medical Institute, Center for Neural Science,
and Courant Institute of Mathematical Sciences, New York University

Correspondence should be addressed to J.W.P. (pillow@gatsby.ucl.ac.uk)

May 23, 2008

Supplementary Materials

To provide more insight into the generalized linear model, we present several figures illustrating details of the model components and their contributions to model response properties. We also discuss several control analyses addressing whether the difference in performance of the full and uncoupled models can be attributed primarily to their difference in capturing the response correlation structure. Finally, we provide connections to several previous results on correlated spiking activity in neural populations, specifically maximum-entropy analyses [1, 2] and pairwise measurements of coding efficiency [3].

Analysis of Model Components

Figure S1 shows the linear bases used to parametrize post-spike filters and coupling filters, respectively. These bases provide a low-dimensional parametrization of the waveforms, allowing for fine temporal

structure near the time of a spike and coarse temporal structure at longer delays. The raised cosine form for these vectors means they sum to a constant and are free of temporal aliasing.

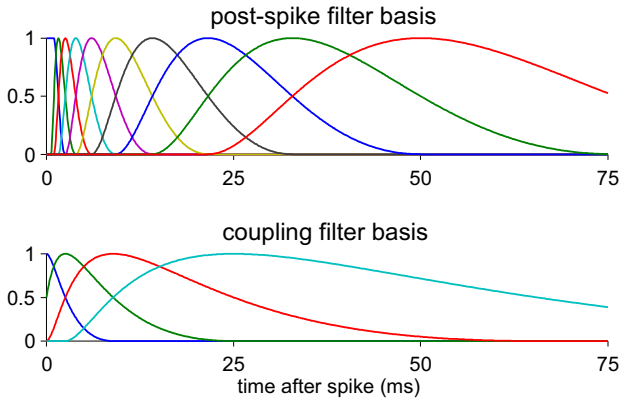

Figure S1: Bases used for representing spike-history and coupling filters. Basis vectors have the form of raised cosine "bumps", so they sum to a constant, and have log-scaling of the time axis, so they can represent fine structure near the time of a spike and coarser structure on longer timescales using a restricted number of parameters (see Methods). Filters were fit as a weighted linear combination of the basis vectors. **Above:** 10-dimensional basis for post-spike filters. **Below:** 4-dimensional basis for coupling filters.

Figure S2 shows that coupling strength falls off strongly with the distance between receptive field (RF) centers, illustrating the fact that coupling effects are primarily restricted to neighboring cells (cf. [2]). Coupling between ON cells is on average much stronger and extends over larger distances than between OFF cells.

Comparison of the fitted parameters obtained for the full and uncoupled models provides insight into the manner in which correlations affect stimulus processing in the retina. Specifically, the spatial extent of the receptive field "surround" mechanism is larger for the uncoupled model than for the full model (fig. S3). This indicates that, when population activity is taken into account, each cell integrates light from an effectively smaller region. In other words, the effect of stimuli far out in the surround can be more parsimoniously explained in terms of population spiking activity. Classical estimates of receptive fields, such as the spike-triggered average, do not resolve such effects. The functional consequence of this change in receptive field structure (fig. S3c) is to make stimulus filters more orthogonal to one another, meaning that stimulus drive is more independent across neurons under the full model.

To examine the relative magnitude of the inputs provided by stimulus and coupling-related model components, we show the net linear input to an example ON cell from the stimulus filter and from coupling filters from ON cell and to OFF cell activity (fig. S4). The sum of these inputs (including

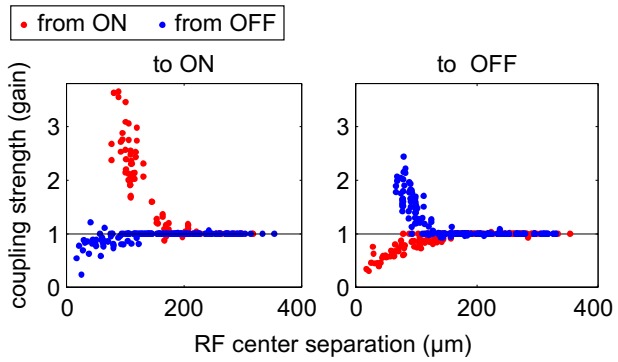

Figure S2: Connectivity summary, showing strength of coupling between cells vs. the distance between their RF centers. Coupling strength is quantified as the exponentiated amplitude of the largest absolute peak of the filter $\exp(f(x_{peak}))$, where $x_{peak} = \arg\max_x |f(x)|$. **Left:** strength of incoming coupling filters to ON cells. **Right:** coupling filters into OFF cells. Points on the line y=1 indicate an absence of coupling (i.e., filters eliminated by the sparse prior). Absence of ON-ON and OFF-OFF points within the range 0–70 μm reflects the minimum spacing between cells within a single mosaic.

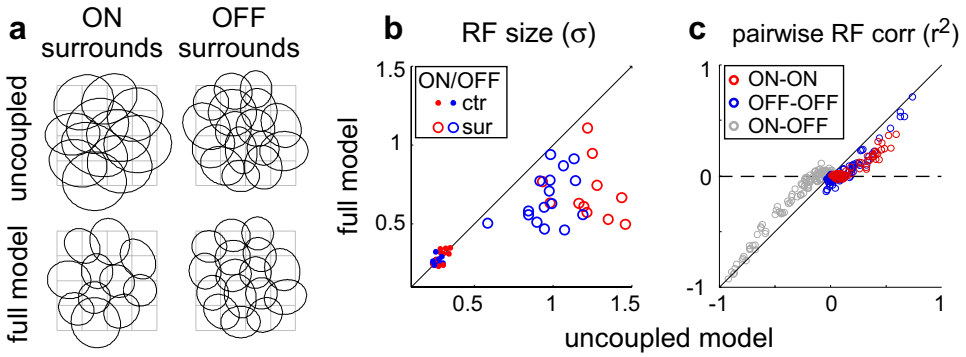

Figure S3: Receptive field changes induced by coupling. **a**, Ellipses show 1 SD contours of a Gaussian fit to the surround mechanism for each cell, for uncoupled (above) and full models (below). **b**, Scatter-plot of RF center and surround sizes; each point represents the estimated RF diameter of a single cell under full and uncoupled models. Surround width is reduced by roughly 70% in the coupled model, while the center width is roughly constant. **c**, Correlation coefficients between all pairs of RFs, under uncoupled (x-axis) and full models (y-axis), showing RFs to be more orthogonal (i.e., closer to zero-correlation) under the full model (cf. [4]).

post-spike filter output) can be loosely interpreted as membrane potential in a "soft-threshold integrate-and-fire" model: the exponential nonlinearity forms a soft threshold whereby the probability of spiking increases as a smooth function of membrane depolarization [5]. The relative importance of network- and stimulus-induced inputs to each cell can be roughly quantified by the amplitude of the fluctuations they induce in the membrane voltage. Across the population, the standard deviation of the total network-induced input is approximately 1/2 the standard deviation of the stimulus-induced input in ON cells, and 1/3 in OFF cells.

For each cell, the model converts linear input into an instantaneous spike rate by means of an expo-

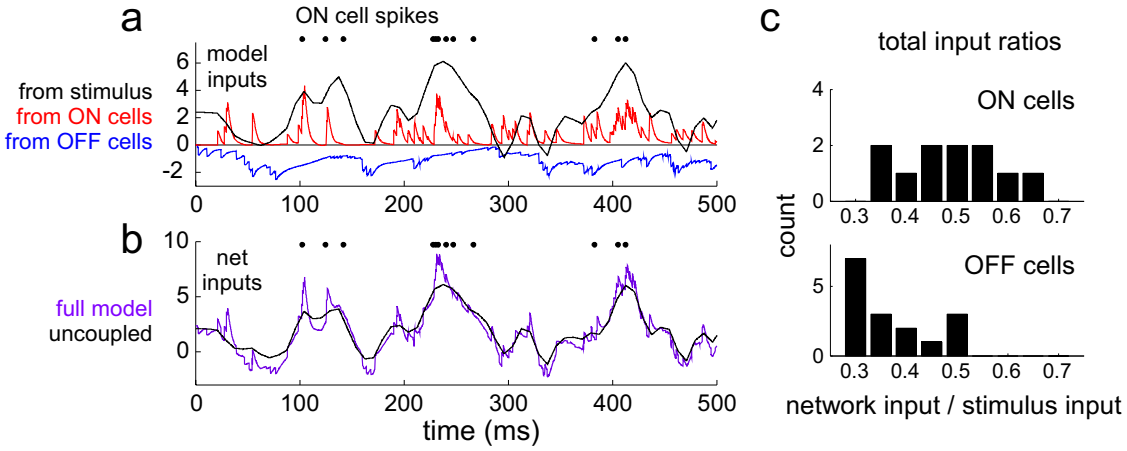

Figure S4: Relative contribution of stimulus and coupling dependent inputs to an example ON cell. **a**, Net linear input from the stimulus, ON cells, and OFF cells during a 500ms stimulus segment on a single trial, with true spike times of the ON cell shown above (black dots). Traces show the output of the stimulus filter (black), and the summed outputs of coupling filters from ON cells (red) and from OFF cells (blue). **b**, Summing the three traces above gives the net external input to the cell (purple trace), plotted alongside the stimulus-dependent input to this cell under the uncoupled model (its only external input). Exponentiating the total input (which also includes input from the post-spike filter, omitted here for visual clarity) gives the instantaneous spike rate of the model. **c**, Histogram showing the relative magnitude of stimulus and population-induced inputs to each cell under the full model. x-axis is the ratio of the standard deviation of each of these inputs (where population input is the sum of ON and OFF inputs). Population-induced input tends to be approximately half as strong as stimulus-induced input in ON cells (above), and about a third as strong in OFF cells (below).

nential nonlinearity. To assess the adequacy of this assumption, we compare an exponential function with a direct "reconstruction" estimate of the nonlinearity, computed using the raw distribution of filter outputs and the observed spike responses (fig. S5) [6]. These reconstructions look reasonably exponential for both uncoupled and full-model parameter settings, though slightly better for ON than OFF cells. For comparison, we also performed a complete re-fitting of the model parameters using output nonlinearities given by a half-wave rectified linear function and by $\log(1+\exp(x))$, which grows linearly for large x and decays like e^x for negative values. These models gave much lower likelihoods for the observed data and exhibited poorer cross-validation performance.

Finally, we re-fit both full and coupled models using a flexible nonlinearity, parametrized as a cubic spline with 8 piecewise polynomial segments. This addition conferred a slight improvement in cross-validation performance (see fig. S7), and it did not elicit a noticeable change in the fitted filters. We

return to this model under Control Analyses, below.

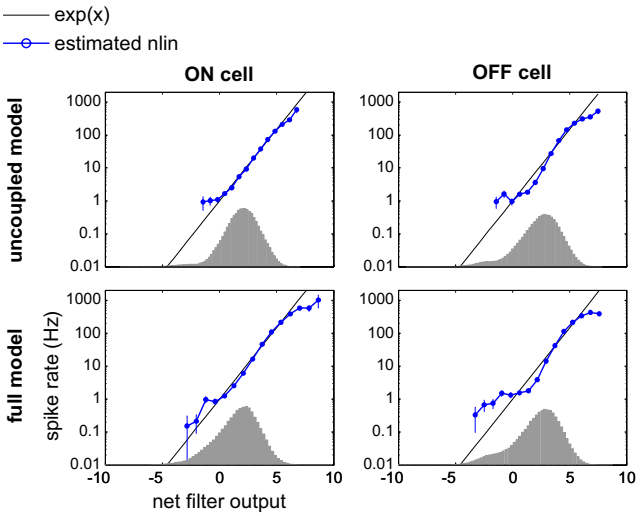

Figure S5: Histogram-based estimates of the nonlinearity transforming linear input to instantaneous spike rate (blue traces), for an example ON (left) and OFF (right) cell, under uncoupled and full models (above and below, respectively) [6]. The gray histograms show the distribution over the net linear input (i.e., the sum of all filter outputs) across time. The nonlinearity represents the probability of observing a spike for each bin in this histogram (with error bars showing ±1 SEM, computed by regarding the spike count in each bin as a binomial random variable). An exponential function (black), the assumed nonlinearity for both models, provides a reasonable approximation to these functions.

Control Analyses

Although they have the same functional form, one might wonder whether the difference in performance of the full and uncoupled models arises from some functional differences other than the full model's ability to incorporate statistical dependencies between neurons. The text of the main article presents several pieces of evidence that the difference is primarily related to correlations. Firstly, the full and uncoupled models predict the PSTH with roughly equal accuracy in all cells, indicating that the models have the same *average* stimulus-response properties. That is, the coupled model does not describe the nonlinear transformation from stimuli to spikes more accurately, once correlations have been averaged out. Secondly, the fact that single-trial predictions correlate more highly with spike trains than the cell's true PSTH (as shown in fig. 3) demonstrates an improvement that *must* be due to correlations: any model that ignores correlations gives the same prediction on every repeat of the stimulus, and therefore cannot do better than the PSTH itself. Thus, capturing the correlation structure clearly play an important role in the coupled model's improved performance.

Nevertheless, we performed several control analyses to determine whether differences in the nonlinear

behavior of the full and uncoupled models might underlie some of our findings. First, we investigated the significance of the difference in model architecture by fitting the parameters of the coupled model to a set of artificial "shuffled data"; we refer to this as the "shuffle-trained model" (fig. S6). This model had the same connectivity and the same number of parameters as the full model, and therefore the same space of possible nonlinear input-output mappings. Shuffled data were generated by using each cell's true spike times in conjunction with simulated spikes from neighboring neurons, elicited by a presentation of the same stimulus to the coupled model.[1] These training data therefore preserved the stimulus dependence of the population response but removed its noise correlation structure. Figure S6 shows the coupling filters obtained for an example ON and OFF cell, compared to those estimated for the full model. Unsurprisingly, the shuffle-trained model exhibited similar PSTH prediction to the full and uncoupled models. It exhibited a negligible improvement in predicting spike times (fig. S6, right), and a $< 1\%$ increase in Bayesian decoding performance (not shown), over the uncoupled model. Therefore, without access to the correlation structure present in simultaneous activity, the full-model architecture behaves like the independent model, and is unable to extract more information from the population responses.

A second group of control analyses was performed using versions of the coupled and uncoupled models with more complicated nonlinear properties (fig. S7). If the full model's advantage over the uncoupled model were due to richer stimulus-processing capabilities, rather than its ability to capture correlations, one would expect to see the advantage disappear once once both models had access to a sufficiently rich class of nonlinearities. We therefore fit the data using several more-complicated "control models", which included:

- a point process model with a flexible (instead of exponential) nonlinearity for each cell, parametrized

[1] We were forced to use simulated data for the shuffled spike-times of neighboring cells due to the fact that we did not have multiple repeats of a long-duration stimulus. We also simulated a full set of simultaneous data from the model to ensure that fitting the full model to this data returned a consistent estimate of the full model parameters. Thus, any stimulus-related features that are capturable by the model were accurately reproduced by this shuffling procedure.

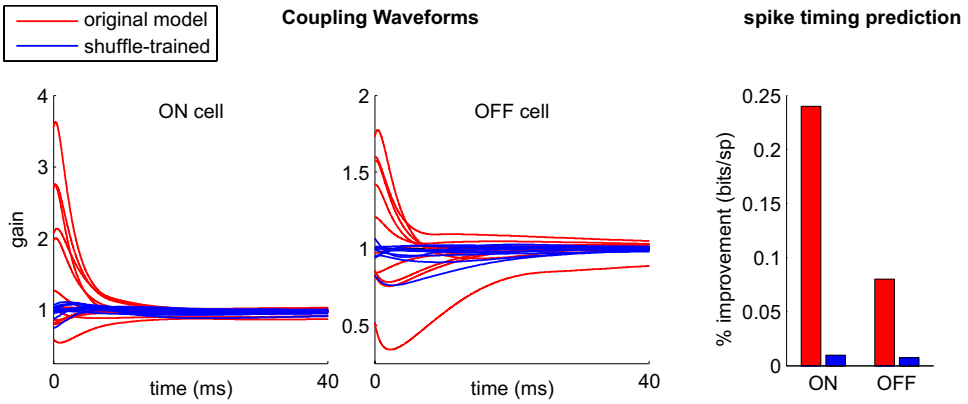

Figure S6: Control analysis #1: coupled model trained using shuffled spike data. The coupled model (assuming the same connectivity as the original coupled model) was fit to artificial "shuffled" data, where the stimulus dependence of the spikes used for fitting was intact, but the correlation structure was removed by shuffling. **Left:** coupling waveforms (blue) obtained for an ON and OFF cell fit to shuffled data, compared with the original waveforms fit to the true (simultaneously recorded) data (red). **Right:** improvement in spike timing prediction of the original full and shuffle-trained models over the uncoupled model.

using cubic splines

- a stochastic, leaky integrate-and-fire model with post-spike current ([7], and compared in [8])
- a point process model with multiple stimulus filters (excitatory and suppressive), giving each cell sensitivity to multiple dimensions of the stimulus space (cf. [9]).

In the third model, the conditional intensity depended on the squared outputs of additional excitatory and suppressive filters, and was given by $\lambda = \lambda_0\Big(\prod_e(1+(k_e\cdot x)^2)\Big)\Big/\Big(\prod_i(1+(k_i\cdot x)^2)\Big)$, where λ_0 is the conditional intensity formulated in the original model, x is the stimulus and $\{k_e\}$ and $\{k_s\}$ are stimulus filters providing quadratic excitatory and suppressive input, respectively. We used two excitatory and two suppressive filters, and fit all parameters (including those governing λ_0) via gradient ascent of the likelihood function. Two versions of each of these control models were fit: one uncoupled, with each cell conditionally independent of the others, and a second with coupling, allowing the model to capture joint dependencies in the response.

For all three control models (fig. S7), the effects of incorporating coupling between cells were directly analogous to the effects in the original model: PSTH prediction was equally accurate with and without

coupling (left panel), while spike-timing prediction (middle) and decoding performance (right) exhibited significant improvements under coupling. Most importantly, the magnitude of the increase in encoding and decoding performance due to coupling was relatively constant across models (i.e., 20% more stimulus information preserved when the population response is decoded under a coupled model). This suggests that the nonlinear stimulus-response properties and response correlation structure make relatively independent contributions to the model's performance. Moreover, these results show that even relatively complicated nonlinear models that ignore correlations do not exceed the performance of the (original) full model. Although not shown here, cross-correlations exhibited by these coupled and uncoupled models are indistinguishable from those of the original full and uncoupled models, indicating that changes in nonlinear stimulus processing do not enhance prediction of the response correlation structure.

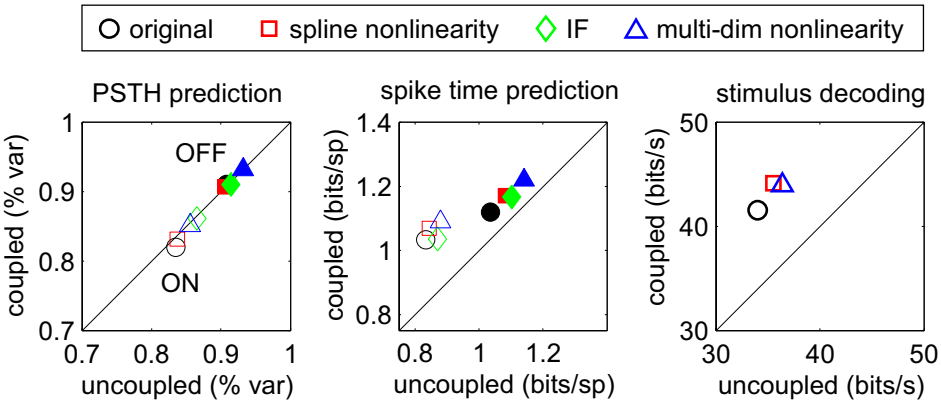

Figure S7: Control analysis #2: comparison of coupled and uncoupled models with more complex nonlinear stimulus processing. Models included: (1) a point process model with a flexible nonlinearity, parametrized using cubic splines (red squares); (2) a stochastic, leaky integrate-and-fire model with post-spike aftercurrent [7]; and (3) a point process model with additional excitatory and suppressive stimulus kernels, allowing multi-dimensional dependence on the stimulus. **Left:** Average percent of the PSTH variance accounted for by coupled and uncoupled variants of each model (open symbols = average over ON cells; filled = average over OFF cells). Prediction improves moderately for the more complicated models, but coupled and uncoupled models exhibit no significant difference. **Middle:** Spike-timing prediction (log-likelihood), showing that coupling provides a similar improvement in single-trial spike prediction across models. **Right:** Bayesian decoding of the population response (stimulus reconstruction) improves slightly for more complex models, but the $\approx 20\%$ improvement conferred by incorporating the response correlation structure is preserved across models (IF model not used due to the computational cost of decoding analysis).

Comparison to maximum-entropy model predictions

Recent work has suggested that the statistics of retinal ganglion cell population responses, both in the presence and absence of stimulation, are well described by a second-order maximum-entropy (or "max-ent") model, which describes the maximum-entropy distribution over binary random variables with fixed mean and covariance [1, 2]. To compare our results with these findings, we fit the second-order max-ent model to the simultaneous responses of a population of seven cells and examined the observed and predicted frequencies of 7-digit binary words, indicating which cells spiked in a 10 ms window (for each cell: 0 = 'no spikes'; 1 = 'one or more spikes'). We compared RGC word frequencies collected during 8 minutes of stimulation with those emitted by the (simulated) max-ent model and coupled/uncoupled generalized linear point process models (fig. S8). The max-ent and coupled point-process models match the observed frequencies with approximately equal accuracy, while the uncoupled-model prediction is noticeably less inaccurate. The coupled model is therefore consistent with predictions of the second-order max-ent model.

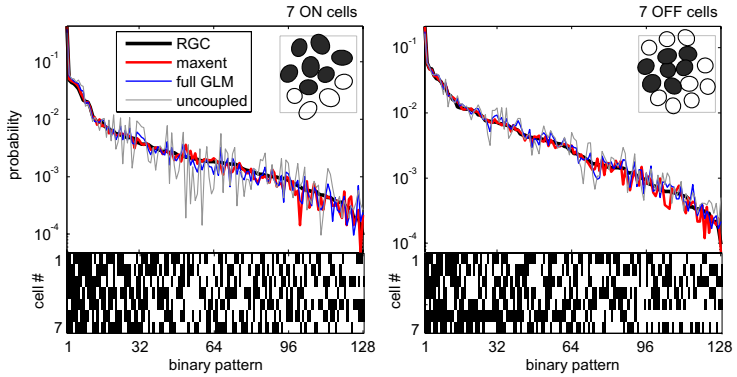

Figure S8: Comparison of 7-neuron binary word frequencies with the predictions of a 2nd-order maximum-entropy model [1, 2]. **Left:** Distribution over binary words, computed using simultaneous responses of 7 ON cells (shown inset) in 10 ms bins. Binary words (shown along x axis) have been sorted by their observed frequency in the actual RGC population response (black trace), and are compared with observed word frequencies from the maximum-entropy (red), coupled (blue) and uncoupled model (gray). **Right:** Analogous plots for a collection of 7 OFF cells. In both cases, the coupled model provides similar accuracy to the max-ent model in predicting word frequencies of RGC population, while the uncoupled model performs substantially less accurately.

However, the generalized linear model makes an important advance over the maximum-entropy model

by incorporating both stimulus dependence and the full time-course of response correlations. Stimulus dependence is especially important, because it is essential for determining whether correlations affect the stimulus-coding properties of the neural population. The max-ent model used here provides a description of the marginal distribution over responses, $P(\mathbf{r})$, while the point-process models describes the *conditional* response distribution, $P(\mathbf{r}|\text{stim})$; only the latter distribution captures the relationship between correlations and stimulus encoding.

More generally, the fact that $P(\mathbf{r})$ can be well described by a second-order maximum entropy model does not imply the same for $P(\mathbf{r}|\text{stim})$. We can illustrate this with a simple example: consider 3 neurons encoding 2 discrete stimuli ('A' and 'B') according to the following rule: if A is presented, an odd number of neurons spike, with each possible spike pattern (100, 010, 001, 111) occurring with equal probability; if stimulus B is presented, an even number of neurons spike, again with all possible patterns (000, 011, 101, 110) equally likely. A second-order max-ent model cannot represent the conditional distributions $P(\mathbf{r}|A)$ and $P(\mathbf{r}|B)$, because third-order correlations are essential to this encoding rule—at least one neuron must have access to both its neighbors' responses when deciding whether or not to spike. However, if A and B occur equally often, then all 8 spike patterns occur with equal probability, so the marginal distribution $P(\mathbf{r})$ is independent, and therefore consistent with a first-order maximum entropy description.

When comparing the second-order maximum-entropy model and the generalized linear point-process model, it is also worth noting that, although both models capture statistical dependencies using pairwise interactions between neurons, responses of the point process model are not necessarily maximum-entropy for any set of constraints on the moments of the response distribution (e.g., 2'nd order correlations). This means that the point-process model does not offer a breakdown of the total entropy by order (i.e., 1st, 2nd, 3rd-order effects), but it might in principle capture higher-order correlations that differ from a second-order maximum entropy prediction.

Connection to pairwise analyses and Poisson spiking model

Previous work on the importance of correlations in neural coding has focused primarily on pairs of neurons, due in part to the large amounts of data needed for information-theoretic (i.e., model-free) analyses of joint coding [3]. To provide a more direct link to these studies, we repeated the decoding analysis shown in fig. 4 using isolated pairs of neurons. Consistent with earlier findings, we observed that preserving the response correlation structure between pairs provides a $\leq 10\%$ increase in information. Our results therefore indicate that the full set of spatio-temporal correlations in a neural population account for more information than is observed for isolated pairs.

Analysis of a full population is critical because pairwise analyses can provide little evidence about the importance of correlations across an entire population. Pairwise measurements are blind to third and higher-order response statistics, meaning there is no limit to the information that could be encoded using higher-order features of the population response. Moreover, second-order pairwise correlations can have radically different influences on the information content of the full population (fig. S9). For an idealized example, consider a population of n neurons, where each neuron contributes I_0 bits independently, and the correlations between any two neurons contribute $0.2I_0$ bits, so that for any two isolated neurons, the percent increase in information due to the correlations is $(0.2I_0)/(2I_0)\times 100 = 10\%$. Now if correlations between any pair of neurons contribute the same information as the correlations between any other pair of neurons (i.e., complete dependence of the pairwise correlations), then the increase due to correlations for the full population is $(0.2I_0)/(nI_0) \times 100 = (20/n)\%$, giving a 0.7% increase for our population of 27 cells. On the other hand, if each set of pairwise correlations contributes independent information, the increase due to correlations for the full population is $\binom{n}{2}0.2I_0)/(nI_0)\times 100 = 10(n-1)\%$, giving a 260% increase for our population (fig. S9). The range of values consistent with pairwise measurements is thus extremely large, and it can increase or decrease with the number of neurons. Determining the importance of correlated spiking across a full population therefore goes significantly beyond the implications of pairwise analyses.

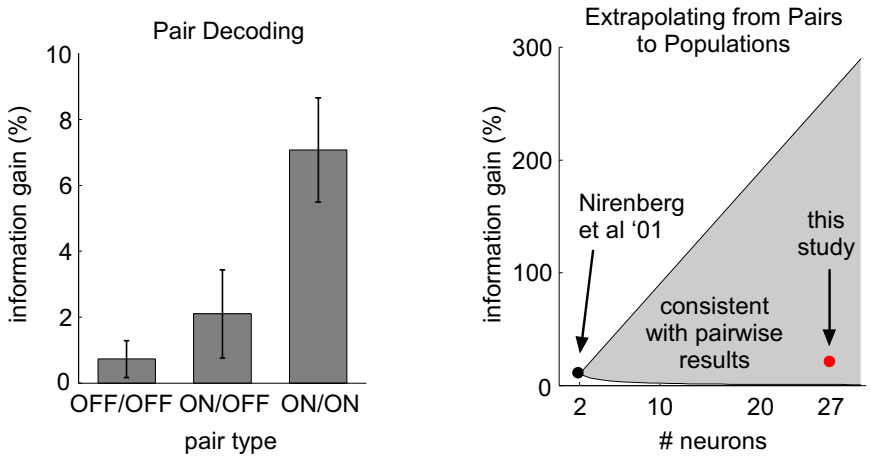

Figure S9: Analysis of pairwise decoding. **Left:** To connect with previous literature [3], we performed the same decoding analysis shown in figure 4 using only two neurons at a time. Each bar represents the mean ($\pm$1SD) percent increase in information for decoding under the coupled model vs. under the independent model, for five OFF pairs, five ON-OFF pairs, and five ON-ON pairs. Consistent with earlier findings, incorporating the correlation structure leads to a $\leq$10% increase in stimulus information. **Right:** Pairwise results alone reveal very little about the importance of correlations in a full population. If including the correlations between two neurons elicits a 10% increase in information (black dot), the 2nd-order correlations in a population of n neurons can elicit an increase between $20/n$% and $10(n-1)$% (gray region). Our results (red dot) pinpoint the value within this range for a modest-sized population.

Lastly, to connect with one of the standard models in the literature, we compared the encoding performance of the generalized linear point-process model with the classic linear-nonlinear-Poisson (LNP) model [10], which lacks both spike-history dependence and coupling between cells (fig. S10). The LNP model predicts both the PSTH (left) and the timing of spikes (right) less accurately than either the full or the uncoupled model, though it still provides more accurate Bayesian stimulus decoding than an optimal linear decoder (fig. 4).

References

[1] E. Schneidman, M. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007–1012, 2006.

[2] J. Shlens, G. Field, J. Gauthier, M. Grivich, D. Petrusca, A. Sher, Litke A. M., and E. J. Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *J Neurosci*, 26:8254–8266, 2006.

[3] S. Nirenberg, S. Carcieri, A. Jacobs, and P. Latham. Retinal ganglion cells act largely as independent encoders. *Nature*, 411:698–701, 2001.

[4] S. H. DeVries. Correlated firing in rabbit retinal ganglion cells. *J. Neurophysiol.*, 81(2):908–920, 1999.

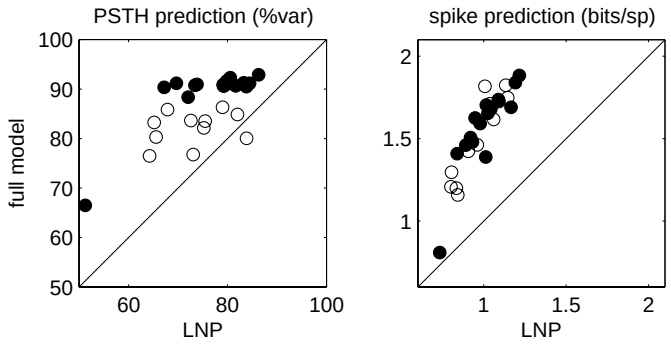

Figure S10: Comparison of predictive power of full model and an inhomogeneous Poisson (LNP) model, which lacks both coupling and post-spike waveforms, and whose output is therefore a Poisson process (c.f. fig. 3c-d in the main text; same conventions apply). The LNP model performs less accurately than the models with spike history effects, as measured both by the PSTH variance accounted for (left) and the log-likelihood of the observed spike data (right) [7].

[5] H. Plesser and W. Gerstner. Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Computation*, 12:367–384, 2000.

[6] E. J. Chichilnisky. A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems*, 12:199–213, 2001.

[7] J. W. Pillow, L. Paninski, V. J. Uzzell, E. P. Simoncelli, and E. J. Chichilnisky. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *The Journal of Neuroscience*, 25:11003–11013, 2005.

[8] L. Paninski, J. W. Pillow, and J. Lewi. Statistical models for neural encoding, decoding, and optimal stimulus design. In P. Cisek, T. Drew, and J. Kalaska, editors, *Computational Neuroscience: Theoretical Insights Into Brain Function*, Progress in Brain Research. Elsevier, 2007.

[9] N. C. Rust, O. Schwartz, J. A. Movshon, and E. P. Simoncelli. Spatiotemporal elements of macaque v1 receptive fields. *Neuron*, 46(6):945–956, 2005.

[10] E. P. Simoncelli, L. Paninski, J. Pillow, and O. Schwartz. Characterization of neural responses with stochastic stimuli. In M. Gazzaniga, editor, *The Cognitive Neurosciences*, pages 327–338. MIT Press, 3rd edition, 2004.