

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

Jonathan W. Pillow¹, Jonathon Shlens², Liam Paninski³, Alexander Sher⁴, Alan M. Litke⁴, E. J. Chichilnisky² & Eero P. Simoncelli⁵

Statistical dependencies in the responses of sensory neurons govern both the amount of stimulus information conveyed and the means by which downstream neurons can extract it. Although a variety of measurements indicate the existence of such dependencies^{1–3}, their origin and importance for neural coding are poorly understood. Here we analyse the functional significance of correlated firing in a complete population of macaque parasol retinal ganglion cells using a model of multi-neuron spike responses^{4,5}. The model, with parameters fit directly to physiological data, simultaneously captures both the stimulus dependence and detailed spatio-temporal correlations in population responses, and provides two insights into the structure of the neural code. First, neural encoding at the population level is less noisy than one would expect from the variability of individual neurons: spike times are more precise, and can be predicted more accurately when the spiking of neighbouring neurons is taken into account. Second, correlations provide additional sensory information: optimal, model-based decoding that exploits the response correlation structure extracts 20% more information about the visual scene than decoding under the assumption of independence, and preserves 40% more visual information than optimal linear decoding⁶. This model-based approach reveals the role of correlated activity in the retinal coding of visual stimuli, and provides a general framework for understanding the importance of correlated activity in populations of neurons.

How does the spiking activity of a neural population represent the sensory environment? The answer depends critically on the structure of neuronal correlations, or the tendency of groups of neurons to fire temporally coordinated spike patterns. The statistics of such patterns have been studied in a variety of brain areas, and their significance in the processing and representation of sensory information has been debated extensively^{2,3,7–13}.

Previous studies have examined visual coding by pairs of neurons¹¹ and the statistics of simultaneous firing patterns in larger neural populations^{14,15}. However, no previous approach has addressed how correlated spiking activity in complete neural populations depends on the pattern of visual stimulation, or has answered the question of how such dependencies affect the encoding of visual stimuli.

Here we introduce a model-based methodology for studying this problem. We describe the encoding of stimuli in the spike trains of a neural population with a generalized linear model (Fig. 1a), a generalization of the well-known linear–nonlinear–Poisson (LNP) cascade model^{4,5,16,17}. In this model, each cell's input is described by a set of linear filters: a stimulus filter, or spatio-temporal receptive field; a post-spike filter, which captures dependencies on spike-train history

(for example, refractoriness, burstiness and adaptation); and a set of coupling filters, which capture dependencies on the recent spiking of other cells. For each neuron, the summed filter responses are exponentiated to obtain an instantaneous spike rate. This is equivalent to exponentiating the filter outputs and then multiplying; the exponentiated post-spike and coupling filters (as plotted in Fig. 1) may therefore be interpreted as spike-induced gain adjustments of the neuron's firing rate.

Although this model is strictly phenomenological, its components can be loosely compared to biophysical mechanisms: the stimulus filter approximates the spatio-temporal integration of light in the outer retina and passive dendritic filtering; the post-spike filter mimics voltage-activated currents following a spike; coupling filters resemble synaptic or electrical interactions between cells (and can mimic the effects of shared input noise); and the exponential non-linearity implements a 'soft threshold', converting membrane potential to instantaneous spike probability. Note that the post-spike and coupling filters, which allow stochastic spiking in one cell to affect subsequent population activity, give rise to shared, non-Poisson variability in the model response.

We fit the model to data recorded *in vitro* from a population of 27 ON and OFF parasol ganglion cells (RGCs) in a small patch of isolated macaque monkey retina, stimulated with 120-Hz spatio-temporal binary white noise. The receptive fields of each of the two cell types formed a complete mosaic covering a small region of visual space (Fig. 1b), indicating that every parasol cell in this region was recorded^{15,18}. Such complete recordings, which have not been achieved elsewhere in the mammalian nervous system, are essential for understanding visual coding in neural populations.

The model contains many parameters that specify the shapes of all filters, but fitting by maximizing likelihood remains highly tractable⁵. A penalty on coupling filters was used to obtain a minimally sufficient set of coupling filters, which yields an estimate of the network's functional connectivity^{19,20}.

Figure 1 shows the estimated filters describing input to example ON and OFF cells. The stimulus filters exhibit centre-surround receptive field organization consistent with previous characterizations of parasol cells. Post-spike filters show the time course of recovery from refractoriness after a spike, and coupling filters show the effects of spikes from nearby cells: for the ON cell (top), spikes in neighbouring ON cells elicit a large, transient excitation (increasing the instantaneous spike rate by a factor of three), whereas spikes in nearby OFF cells elicit suppression. These effects are reversed in the OFF cell, which is excited/suppressed by spikes in neighbouring OFF/ON cells. Both populations exhibit approximate nearest-neighbour connectivity, with coupling strength falling as a function of distance

¹Gatsby Computational Neuroscience Unit, UCL, 17 Queen Square, London WC1N 3AR, UK. ²The Salk Institute, 10010 North Torrey Pines Road, San Diego, California 92037, USA. ³Department of Statistics and Center for Theoretical Neuroscience, Columbia University, 1255 Amsterdam Avenue, New York, New York 10027, USA. ⁴Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, California 95064, USA. ⁵Howard Hughes Medical Institute, Center for Neural Science, and Courant Institute of Mathematical Sciences, New York University, 4 Washington Place, Room 809, New York, New York 10003, USA.

between receptive field centres¹⁵. We found that fitted stimulus filters have smaller surrounds than the spike-triggered average, indicating that a portion of the classical surround can be explained by interactions between cells²¹ (see Supplementary Information).

To assess accuracy in capturing the statistical dependencies in population responses, we compared the pairwise cross-correlation function (CCF) of RGCs and simulated model spike trains (Fig. 2). For nearby ON–ON and OFF–OFF pairs, the CCF exhibits a sharp peak at zero, indicating the prevalence of synchronous spikes; however, for ON–OFF pairs, a trough at zero indicates an absence of synchrony. For all 351 possible pairings, the model accurately reproduces the CCF (Fig. 2a–c, e, f).

To examine whether inter-neuronal coupling was necessary to capture the response correlation structure, we re-fitted the model without coupling filters (that is, so that each cell's response depends only on the stimulus and its own spike-train history). This 'uncoupled model' assumes that cells encode the stimulus independently, although correlations may still arise from the overlap of stimulus filters. However, the uncoupled model fails to reproduce the sharp CCF peaks observed in the data. These peaks are also absent from CCFs computed on trial-shuffled data, indicating that fast-timescale correlations are not stimulus-induced and therefore cannot be captured by any independent encoding model.

Higher-order statistical dependencies were considered by inspecting correlations in three-neuron groups: triplet CCFs show the spike rate of one cell as a function of the relative time to spikes in two other cells (Fig. 2e–g)¹⁵. For adjacent neurons of the same type, triplet CCFs have substantial peaks at zero ('triplet synchrony'), which are well matched by the full model.

Although the full and uncoupled models differ substantially in their statistical dependencies, the two models predict average light responses in individual cells with nearly identical accuracy, capturing 80–95% of the variance in the peri-stimulus time histogram (PSTH) in 26 out of 27 cells (Fig. 3a–c). Both models therefore accurately describe average single-cell responses to new stimuli. However, the full model achieves higher accuracy, predicting multi-neuronal spike responses on a single trial ($8 \pm 3\%$ more bits per spike, Fig. 3d). This discrepancy can be explained by the fact that noise is shared across

neurons. Shared variability means that population activity carries information about a single cell's response (owing to coupling between cells) beyond that provided by the stimulus alone. Individual neurons therefore appear less noisy when conditioned on spiking activity in the rest of the population than they appear in raster plots.

We measured the effect of correlations on single-trial, single-cell spike-train prediction by using the model to draw samples of a single cell's response given both the stimulus and the spiking activity in the rest of the population on a single trial (Fig. 3e, f). Averaging the resulting raster plot gives a prediction of the cell's single-trial spike rate, or 'population-conditioned' PSTH for a single trial. We compared these predictions with the cell's true spike times (binned at 2 ms) across all trials and found that on nearly every trial, the model-based prediction is more highly correlated with the observed spikes than the neuron's full PSTH (Fig. 3g). Note that the full PSTH achieves the highest correlation possible for any trial-independent prediction. Thus, by exploiting the correlation structure, the coupled model predicts single-neuron spike times more accurately than any independent encoding model.

Although the full model accurately captures dependencies in the activity of RGCs, it is not obvious *a priori* whether these dependencies affect the amount of sensory information conveyed by RGC responses. In principle, the correlation structure could be necessary to predict the responses, but not to extract the stimulus information that the responses carry¹³. To examine this issue directly, we used the full and uncoupled models to perform Bayesian decoding of the population response (Fig. 4a), which optimally reconstructs stimuli given an accurate description of the encoding process. For comparison, we also performed Bayesian decoding under a Poisson (that is, LNP) model and optimal linear decoding⁶.

Each decoding method was used to estimate short (150-ms) segments of the stimulus given all relevant spike times from the full population (Fig. 4b). Bayesian decoding under the coupled model recovers 20% more information than Bayesian decoding under the uncoupled model, indicating that knowledge of the correlation structure is critical for extracting all sensory information contained in the population response. This improvement was invariant to enhancements of the

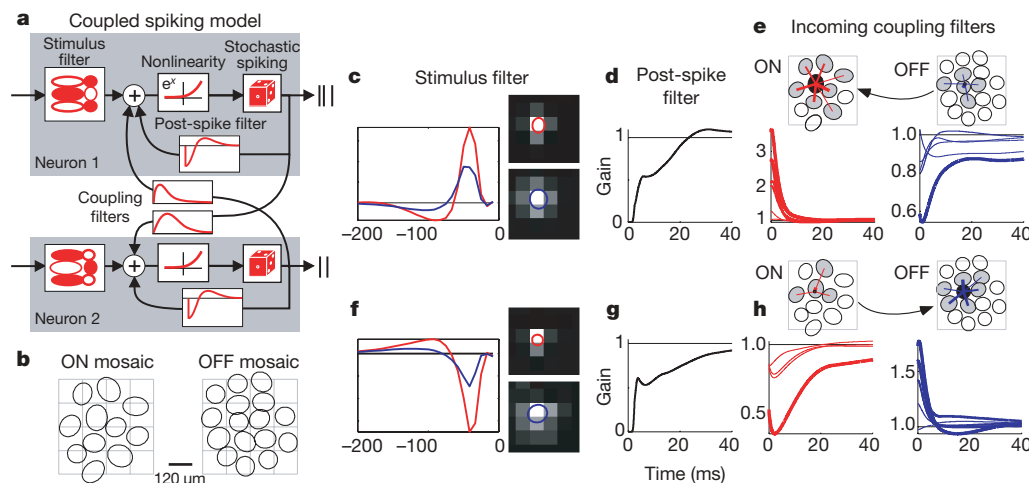


Figure 1 | Multi-neuron encoding model and fitted parameters. **a**, Model schematic for two coupled neurons: each neuron has a stimulus filter, a post-spike filter and coupling filters that capture dependencies on spiking in other neurons. Summed filter output passes through an exponential nonlinearity to produce the instantaneous spike rate. **b**, Mosaics of 11 ON and 16 OFF retinal ganglion cell receptive fields, tiling a small region of visual space. Ellipses represent 1 s.d. of a Gaussian fit to each receptive field centre; the square grid indicates stimulus pixels. **c–e**, Parameters for an example ON cell. **c**, Temporal and spatial components of centre (red) and surround (blue) filter components, the difference of which is the full stimulus filter.

d, Exponentiated post-spike filter, which may be interpreted as multiplying the spike rate after a spike at time zero. It produces a brief refractory period and gradual recovery (with a slight overshoot). **e**, Connectivity and coupling filters from other cells in the population. The black filled ellipse is this cell's RF centre, and blue and red lines show connections from neighbouring OFF and ON cells, respectively (line thickness indicates coupling strength). Below, exponentiated coupling filters show the multiplicative effect on this cell's spike rate after a spike in a neighbouring cell. **f–h**, Analogous plots for an example OFF cell.

model's stimulus filters and nonlinearities (see Supplementary Information), indicating that the difference in performance arises specifically from the coupled model's ability to incorporate the correlation structure. Our results also show that spike history is relevant for decoding (a Poisson model preserves 6% less information than the uncoupled model²²) and that restricting to a linear decoder further reduces the information that can be recovered from RGC responses.

Decoding analysis can also be used to examine the coding fidelity of specific stimulus features. As a simple illustration, we examined the temporal frequency spectrum of reconstructed stimuli and found that the response correlation structure is most important for decoding those stimulus frequencies (6–20 Hz) that are encoded with highest fidelity (Fig. 4c).

These results demonstrate that the responses of a population of retinal ganglion cells are well described by a generalized linear model, and that correlations in the response can be exploited to recover 20% more visual information than if responses are regarded as independent given the stimulus. In contrast, previous studies have reported this information gain to be less than 10% for pairs of neurons^{9,12}. However, pairwise analyses provide little evidence about the importance of correlations across an entire population. Second-order correlations between pairs of neurons could give rise to either much larger (scaling with the number of neurons *n*) or much smaller (falling as 1/*n*) gains for a full population (see Supplementary Information). To compare more directly with previous findings,

we performed Bayesian decoding using isolated pairs of neurons from the same population; we found a $\leq 10\%$ gain in sensory information when correlations were included (see Supplementary Information). This is consistent with previous findings, and shows that the information gain for a complete population is larger than that observed for pairs. We also compared the model to a pairwise maximum-entropy model, which has recently been shown to capture the instantaneous spiking statistics of groups of retinal ganglion cells^{14,15}. The coupled model exhibits similar accuracy in capturing these statistics, but has the advantage that it accounts for the temporal correlation structure and stimulus dependence of responses, which are essential for assessing the effect of correlations on sensory coding.

Although it provides an accurate functional description of correlated spike responses, the generalized linear model does not reveal the biophysical mechanisms underlying the statistical dependencies between neurons: coupling does not necessarily imply anatomical connections between cells, but could (for example) reflect dependencies due to shared input noise¹. The model also lacks several mechanisms known to exist in retinal ganglion cells (for example, contrast gain-control²³), which may be required for characterizing responses to a wider variety of stimuli. One additional caveat is that Bayesian decoding provides a tool for measuring the sensory information available in the population response, but it does not reveal whether the brain makes use of this information.

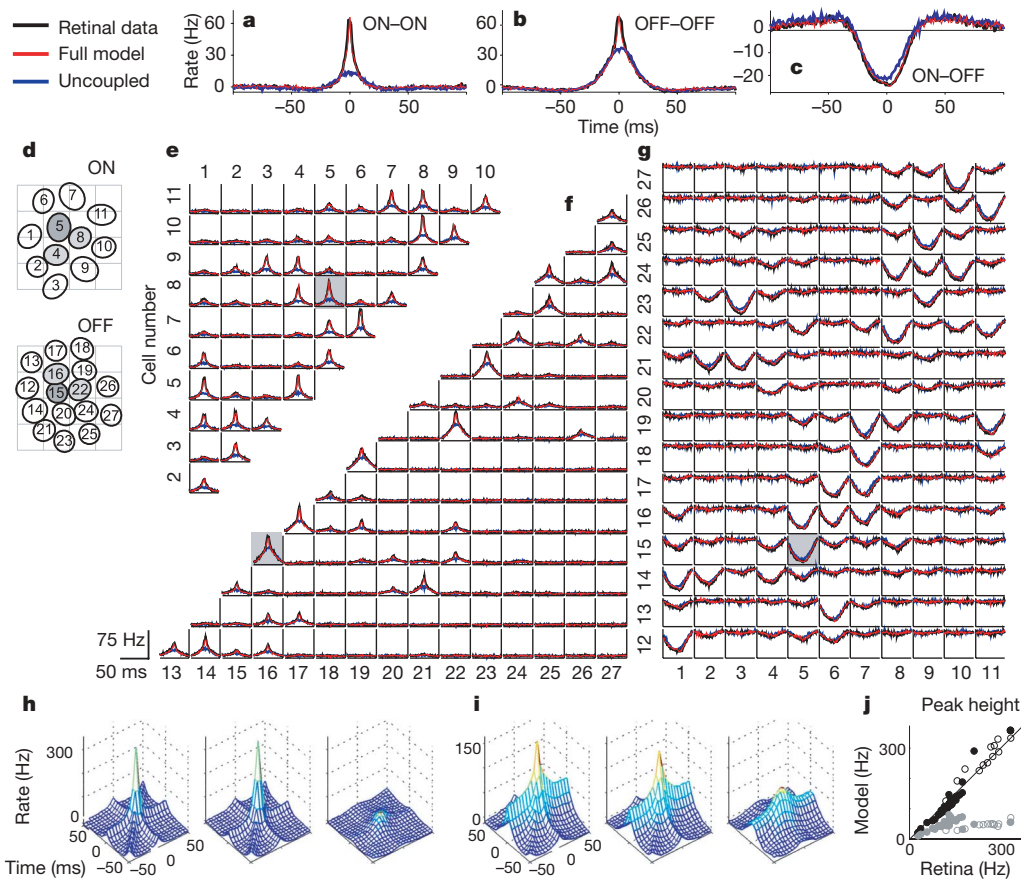


Figure 2 | Analysis of response correlations. **a–c**, Example CCFs of retinal responses, and simulated responses of the full and uncoupled models, for two ON cells (**a**), two OFF cells (**b**) and an ON–OFF pair (**c**). The baseline is subtracted so that units are in spikes per s above (or below) the cell's mean rate. **d**, Receptive field mosaic overlaid with arbitrary labels. Dark grey indicates cells shown in Fig. 1; light grey indicates cells used for triple correlations (**h**, **i**). **e**, CCFs between all ON pairs, where the *i*, *j*th plot shows the CCF between cell *i* and cell *j*. The grey box indicates the CCF plotted in

a, **f**, **g**, CCFs between all OFF–OFF pairs (**f**) and all ON–OFF pairs (**g**; abscissa height is 30 Hz). **h**, Third-order (triplet) CCF between three adjacent ON cells, showing the instantaneous spike rate of cell 5 as a function of the relative spike time in cells 4 and 8 (left, RGCs; middle, full model; right, uncoupled model). **i**, Analogous triplet CCF for OFF cells 15, 16 and 22. **j**, Comparison of the triplet CCF peak in RGC and model responses (full model, black; uncoupled, grey) for randomly selected triplets of adjacent ON (open) and OFF (filled) cells.

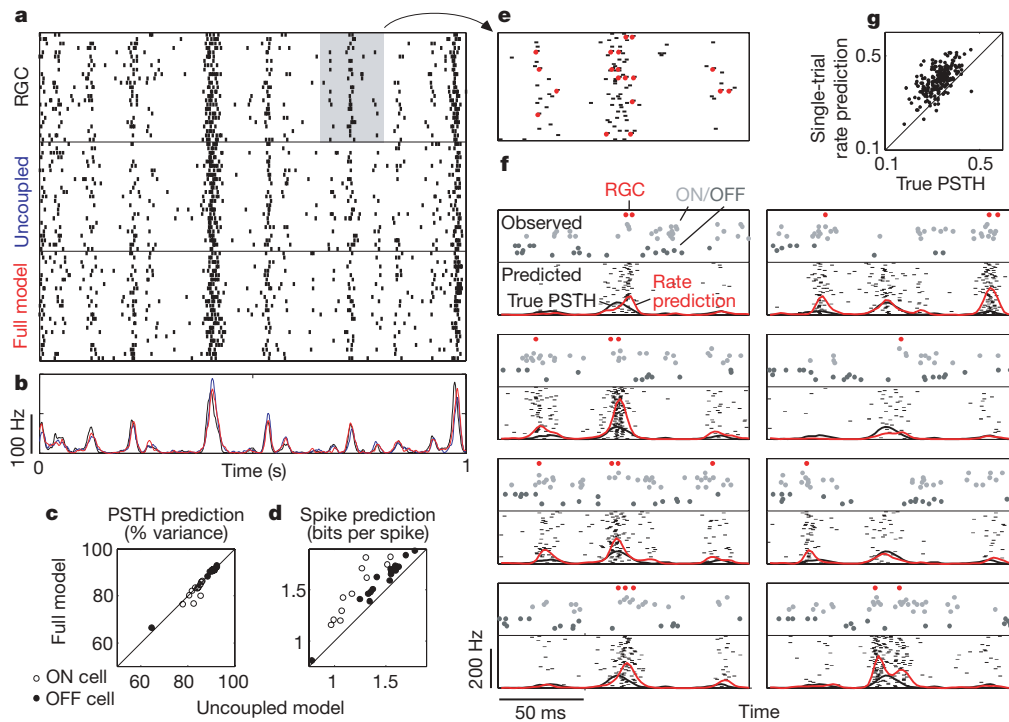


Figure 3 | Spike-train prediction comparison. **a**, Raster of responses of an ON RGC to 25 repeats of a novel 1-s stimulus (top), and responses of uncoupled (middle) and full (bottom) models to the same stimulus. **b**, PSTH of the RGC (black), uncoupled (blue) and coupled (red) model; both models account for ~84% of the variance of the true PSTH. **c**, PSTH prediction by full and uncoupled models, showing that coupling confers no advantage in predicting average responses. **d**, Log-likelihood of novel RGC spike responses under full and uncoupled models; the full model provides 8% more information about novel spike trains. **e**, Magnified 150-ms portion of RGC raster and PSTH (grey box in **a**). Red dots highlight RGC spike times on selected individual trials, which are replotted in **f**. **f**, Single-trial spike-train prediction using the coupled model. The top half of each plot shows the

population activity on a single trial: true spike times of the cell (red dots), coupled ON cells (light grey dots) and coupled OFF cells (dark grey dots; each line in the raster shows the spike times of a different cell). The bottom half of each plot shows a raster of 50 predicted responses of the cell in question, using both the stimulus and coupled responses (shown above) to predict spike trains. The red trace shows the single-trial rate prediction (population-conditioned PSTH), compared with true PSTH of the cell (black trace, identical in all plots). **g**, Correlation coefficient of true spike trains with the PSTH (ordinate) and with population-conditioned predictions (abscissa); the full model predicts single-trial responses with higher accuracy than the true PSTH.

Physiological interpretations of the model and mechanisms for neural read-out of sensory information in higher brain areas are thus important directions for future research.

Nevertheless, the generalized linear model offers a concise, computationally tractable description of the population encoding process, and provides the first generative description of the space–time dependencies in stimulus-induced population activity. It allows us to quantify the relative contributions of stimulus, spike history and

network interactions to the encoding and decoding of visual stimuli, and clarifies the relationship between single-cell and population variability. More generally, the model can be used to assess which features of the visual environment are encoded with highest and lowest fidelity, and to determine how the structure of the neural code constrains perceptual capabilities. We expect this framework to extend to other brain areas, and to have an important role in revealing the information processing capabilities of spiking neural populations^{4,19,24,25}.

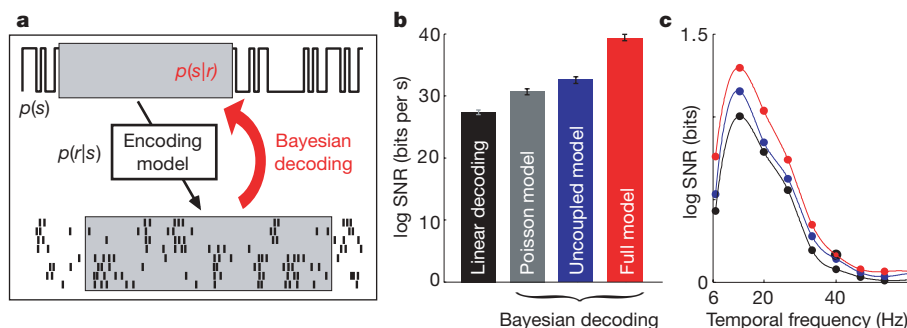


Figure 4 | Decoding performance comparison. **a**, Shown is a Bayesian decoding schematic: to estimate an unknown stimulus segment from a set of observed spike times (highlighted in boxes), the stimulus prior distribution $p(s)$ is multiplied by the model-defined likelihood $p(r|s)$ to obtain the posterior $p(s|r)$. The posterior mean is the Bayes' least-squares stimulus estimate. **b**, Log of the SNR for linear decoding, as well as for Bayesian decoding under the Poisson, uncoupled and full models⁶. The full model

preserves 20% more information than the uncoupled model, which indicates that there is additional sensory information available from the population response when correlations are taken into account. Error bars show 95% confidence intervals based on 2,000 bootstrap resamplings of 3,000 decoded stimulus segments. **c**, Log SNR decomposed as a function of temporal frequency for various decoding methods (Poisson omitted for clarity).

METHODS SUMMARY

Data. Multi-electrode extracellular recordings were obtained *in vitro* from a segment of isolated, peripheral macaque monkey (*Macaca mulatta*) retina, and analysis was restricted to two cell types (ON and OFF parasol)^{15,26,27}. A standard spike-sorting procedure, followed by a specialized statistical method for detecting simultaneous spikes, was used to sort spikes (see ref. 28). The retina was stimulated with a photopic, achromatic, optically reduced spatio-temporal binary white noise stimulus refreshing at 120 Hz, with a root-mean-square contrast of 96%.

Fitting. Model parameters were fitted to 7 min of spike responses to a non-repeating stimulus. Each cell's parameters consisted of a stimulus filter (parametrized as a rank-2 matrix), a spike-history filter, a set of incoming coupling filters and a constant. Temporal filters were represented in a basis of cosine 'bumps'²². Parameters for the uncoupled and Poisson (LNP) models were fitted independently. Parameters were fitted by penalized maximum likelihood^{4,5}, with an L1 penalty on the vector length of coupling filters to eliminate unnecessary connections.

Encoding. Spike prediction was cross-validated using the log-likelihood of 5 min of novel spiking data (scaled to units of bits per s). Repeat rasters were obtained using 200 presentations of a novel 10-s stimulus. Population-conditional rasters were obtained from the coupled model by sampling the model-defined probability distribution over the neuron's response given the stimulus and surrounding-population activity on a single trial²⁹.

Decoding. Population responses were decoded using the Bayes' least-squares estimator (posterior mean) to reconstruct 18-sample single-pixel stimulus segments (cross-validation data). Linear decoding was performed using the optimal linear estimator⁶. Decoding performance was quantified using the log signal-to-noise ratio (SNR) of each technique, which gives an estimate of mutual information. Breakdown by temporal frequency was obtained by computing the Fourier power spectra of the stimuli and residuals and then computing log SNR.

Full Methods and any associated references are available in the online version of the paper at www.nature.com/nature.

Received 16 July 2007; accepted 5 June 2008.

Published online 23 July 2008.

- Mastrorade, D. N. Correlated firing of retinal ganglion cells. *Trends Neurosci.* **12**, 75–80 (1989).
- Meister, M., Lagnado, L. & Baylor, D. A. Concerted signaling by retinal ganglion cells. *Science* **270**, 1207–1210 (1995).
- Shadlen, M. & Newsome, W. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P. & Brown, E. N. A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *J. Neurophysiol.* **93**, 1074–1089 (2004).
- Paninski, L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network Comp. Neural Syst.* **15**, 243–262 (2004).
- Warland, D., Reinagel, P. & Meister, M. Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* **78**, 2336–2350 (1997).
- Dan, Y., Alonso, J. M., Usrey, W. M. & Reid, R. C. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neurosci.* **1**, 501–507 (1998).
- Panzeri, S., Golledge, H., Zheng, F., Tovee, M. P. & Young, M. J. Objective assessment of the functional role of spike train correlations using information measures. *Vis. Cogn.* **8**, 531–547 (2001).
- Nirenberg, S., Carcieri, S., Jacobs, A. & Latham, P. Retinal ganglion cells act largely as independent encoders. *Nature* **411**, 698–701 (2001).
- Schneidman, E., Bialek, W. & Berry, M. J. Synergy, redundancy, and independence in population codes. *J. Neurosci.* **21**, 11539–11553 (2003).
- Nirenberg, S. & Latham, P. E. Decoding neuronal spike trains: how important are correlations? *Proc. Natl Acad. Sci. USA* **100**, 7348–7353 (2003).
- Averbeck, B. B. & Lee, D. Coding and transmission of information by neural ensembles. *Trends Neurosci.* **27**, 225–230 (2004).
- Latham, P. & Nirenberg, S. Synergy, redundancy, and independence in population codes, revisited. *J. Neurosci.* **25**, 5195–5206 (2005).
- Schneidman, E., Berry, M., Segev, R. & Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).
- Shlens, J. et al. The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.* **26**, 8254–8266 (2006).
- Plesser, H. & Gerstner, W. Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Comput.* **12**, 367–384 (2000).
- Simoncelli, E. P., Paninski, L., Pillow, J. & Schwartz, O. in *The Cognitive Neurosciences* 3rd edn (ed. Gazzaniga, M.) 327–338 (MIT, 2004).
- Frechette, E. S. et al. Fidelity of the ensemble code for visual motion in primate retina. *J. Neurophysiol.* **94**, 119–135 (2005).
- Okatan, M., Wilson, M. & Brown, E. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput.* **17**, 1927–1961 (2005).
- Rigat, F., de Gunst, M. & van Pelt, J. Bayesian modelling and analysis of spatio-temporal neuronal networks. *Bayes. Anal.* **1**, 733–764 (2006).
- DeVries, S. H. Correlated firing in rabbit retinal ganglion cells. *J. Neurophysiol.* **81**, 908–920 (1999).
- Pillow, J. W., Paninski, L., Uzzell, V. J., Simoncelli, E. P. & Chichilnisky, E. J. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *J. Neurosci.* **25**, 11003–11013 (2005).
- Shapley, R. M. & Victor, J. D. The effect of contrast on the transfer properties of cat retinal ganglion cells. *J. Physiol.* **285**, 275–298 (1978).
- Harris, K., Csicsvari, J., Hirase, H., Dragoi, G. & Buzsaki, G. Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).
- Paninski, L., Fellows, M., Shoham, S., Hatsopoulos, N. & Donoghue, J. Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J. Neurosci.* **24**, 8551–8561 (2004).
- Litke, A. M. et al. What does the eye tell the brain? Development of a system for the large scale recording of retinal output activity. *IEEE Trans. Nucl. Sci.* **51**, 1434–1440 (2004).
- Watanabe, M. & Rodieck, R. W. Parasol and midget ganglion cells of the primate retina. *J. Comp. Neurol.* **289**, 434–454 (1989).
- Segev, R., Goodhouse, J., Puchalla, J. & Berry, M. J. Recording spikes from a large fraction of the ganglion cells in a retinal patch. *Nature Neurosci.* **7**, 1155–1162 (2004).
- Pillow, J. W. & Latham, P. in *Advances in Neural Information Processing Systems 20* (eds Platt, J. C., Koller, D., Singer, Y. & Roweis, S.) 1161–1168 (MIT, 2008).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We thank M. Bethge, C. Brody, D. Butts, P. Latham, M. Lengyel, S. Nirenberg and R. Sussman for comments and discussions; G. Field, M. Greschner, J. Gauthier and C. Hulse for experimental assistance; M. I. Grivich, D. Petrusca, W. Dabrowski, A. Grillo, P. Grybos, P. Hottowy and S. Kachiguine for technical development; H. Fox, M. Taffe, E. Callaway and K. Osborn for providing access to retinas; and S. Barry for machining. Funding was provided a Royal Society USA/Canada Research Fellowship (J.W.P.); NSF IGERT DGE-03345 (J.S.); NEI grant EY018003 (E.J.C., L.P. and E.P.S.); Gatsby Foundation Pilot Grant (L.P.); Burroughs Wellcome Fund Career Award at the Scientific Interface (A.S.); US National Science Foundation grant PHY-0417175 (A.M.L.); McKnight Foundation (A.M.L. and E.J.C.); and HHMI (J.W.P., L.P. and E.P.S.).

Author Information Reprints and permissions information is available at www.nature.com/reprints. Correspondence and requests for materials should be addressed to J.W.P. (pillow@gatsby.ucl.ac.uk).

METHODS

Recording. Multi-electrode extracellular recordings were obtained *in vitro* from a segment of isolated, peripheral macaque monkey retina, using preparation and recording methods described previously^{15,26}. Analysis was restricted to two physiologically defined classes of cells; on the basis of light response properties and density, these were identified as ON and OFF parasol cells²⁷. The cells shown were recorded in a square region of retina covered by 76 electrodes. A standard clustering-based spike-sorting procedure (see refs 15, 26) was used to estimate the number of units, and least-squares regression of the estimated spike times against multi-electrode voltage signal was used to estimate multi-electrode spike waveforms for each unit. Although this approach correctly and efficiently identifies isolated spikes, when two cells fire within a 1–2-ms window, the clustering approach can fail to identify the presence of both spikes. We solved this problem by using estimates of the elementary waveforms to detect the superposition of spikes. We performed maximum *a posteriori* estimation under the model that the multi-electrode voltage signal was the linear superposition of Gaussian white noise and the spike trains convolved with their associated spike waveforms, with a sparse (exponential) prior distribution on the spike trains. This corresponds to a tractable quadratic optimization problem under linear inequality constraints, which can be solved efficiently using existing methods. The real-valued solution vector was then binarized by greedily inserting spikes whenever the reduction in mean-squared error between predicted and actual voltage exceeded a threshold²⁸. This procedure correctly identified simultaneous spikes in simulated data sets and corrected obvious cross-correlation artefacts appearing in real data sorted with standard clustering techniques.

Stimuli. The retina was stimulated with a photopic, achromatic image of a cathode ray tube display, refreshing at 120 Hz. The stimulus was a spatio-temporal pseudo-random binary sequence, where the intensity of each pixel was drawn independently from one of two values on each frame. The stimulus pixel size was $120 \times 120 \mu\text{m}$ on the retina, and contrast (standard deviation divided by mean) was 96%.

Fitting. Model parameters were fitted by maximizing likelihood⁵ using 7 min of spiking data recorded during presentation of a non-repeating stimulus. The parameters for each cell consisted of a stimulus filter \mathbf{k} , a spike-history filter \mathbf{h} , a set of incoming coupling filters $\{\mathbf{l}_i\}$ and a constant (specifying the log of the baseline firing rate) μ . The filter \mathbf{k} was a 750-dimensional vector (5×5 spatial pixels \times 30 time bins), parametrized using a lower-dimensional representation as a rank-2 matrix: $k(x, y, \tau) = k_{s,1}(x, y)k_{t,1}(\tau) + k_{s,2}(x, y)k_{t,2}(\tau)$, with $k_{s,i}(x, y)$ denoting a spatial filter (25 parameters) and $k_{t,i}(\tau)$ a temporal filter (10 parameters), giving $2 \times 35 = 70$ parameters. A rank-3 representation did not improve performance. These filters closely resembled a time-varying difference-of-Gaussians³⁰; spatial filters were well-approximated (in a least-squares sense) by Gaussians, which were used to plot spatial ellipses shown in Fig. 1 and to summarize receptive field properties (Supplementary Figs 2 and 3). Gaussians fit to receptive field centres and surrounds had average standard deviations of 0.25 pixels and 0.7 pixels (1.0 pixels for the uncoupled model), respectively. Temporal filters \mathbf{h} and $\{\mathbf{l}_i\}$ and the temporal components of \mathbf{k} were represented using a basis of raised cosine ‘bumps’ of the form $\mathbf{b}_j(t) = (1/2)\cos(a\log[t+c] - \phi_j) + (1/2)$ for t such that $a\log(t+c) \in [\phi_j - \pi, \phi_j + \pi]$ and 0 elsewhere, with constants a and c set by hand to watch the structure observed in auto- and cross-correlation functions, and $\pi/2$ spacing between the ϕ_j (see Supplementary Information). This basis allows for the representation of fine temporal structure near the time of a spike and coarser/smoother dependency at later times (see ref. 22). The \mathbf{h} filter was represented with ten such basis vectors, and the \mathbf{l}_i coupling filters were represented with four. The ‘uncoupled model’ was fitted independently without coupling filters $\{\mathbf{l}_i\}$, and the inhomogeneous Poisson model (Fig. 4) was fitted without $\{\mathbf{l}_i\}$ or \mathbf{h} .

Conditional intensity (spike rate) is given by $\lambda(t) = \exp(\mathbf{k} \cdot \mathbf{x} + \mathbf{h} \cdot \mathbf{y} + (\sum_i \mathbf{l}_i \cdot \mathbf{y}_i) + \mu)$, where \mathbf{x} is the stimulus, \mathbf{y} the cell’s own spike-train history, μ is the cell’s baseline log-firing rate, and $\{\mathbf{y}_i\}$ the spike-train histories of other cells at time t . The population log-likelihood is the sum over single-cell

log-likelihoods, each given by $L = \sum \log \lambda(t_{sp}) - \lambda(t) dt$, where t_{sp} denotes the set of spike times and the integral is taken over the length of the experiment^{4,5}. We added a penalty of the form $-\alpha \int |\sum_i \mathbf{l}_i(t)|^{1/2} dt$ to eliminate unnecessary coupling filters (using a constrained Newton–Raphson algorithm to maximize the penalized log-likelihood), which regularizes and prevents overfitting. The regularization parameter α was selected by means of cross-validation on a novel 5-min data set, but results were robust with respect to both α and the choice of basis. (This reduced the number of coupling filters from 702 to 243 and recovered a roughly pairwise-adjacent structure; see Supplementary Information.)

Correlations. Spike responses of full and uncoupled models were simulated with the same 20-min stimulus (144,000 samples) presented experimentally. Pairwise cross-correlations were computed in 1-ms bins, according to $C(\tau) = [\langle y_1(t)y_2(t+\tau) \rangle - \langle y_1(t) \rangle \langle y_2(t) \rangle] / (\langle y_2(t) \rangle dt)$, where $y_1(t)$ denotes the spike response of the first neuron in bins of width dt , and $\langle \cdot \rangle$ denotes averaging over t . Triplet correlations were computed in 5-ms bins according to $C(\tau_1, \tau_2) = [\langle y_1(t)y_2(t+\tau_1)y_3(t+\tau_2) \rangle - \langle y_1(t) \rangle \langle y_2(t) \rangle \langle y_3(t) \rangle] / (\langle y_2(t) \rangle \langle y_3(t) \rangle dt)$.

Encoding. Spike-train prediction was validated using the log-likelihood of novel spike trains under both models, computed on 5 min of data not used for fitting or setting α . The difference of log-likelihood under the model and log-likelihood under a homogeneous Poisson process, $\sum \log \lambda(t_{sp}) - \int \lambda(t) dt$ (where $\bar{\lambda} = n_{sp}/T$ is the mean spike rate), divided by n_{sp} , gives prediction accuracy in bits per spike for each cell²⁵. Repeat rasters were obtained using 200 presentations of a novel 10-s stimulus, and the time-varying average response (PSTH) was computed in 1-ms bins, smoothed with a Gaussian kernel of width $\sigma = 2$ ms. Conditional rasters were obtained from the coupled model by holding the responses of all but one neuron fixed, and sampling from the model-induced probability distribution on the remaining neuron’s response. Samples were obtained by the Metropolis–Hastings algorithm, with spike ‘proposals’ drawn from a point process model as described in ref. 29. We kept only every 100th output sample of the algorithm to ensure independent samples.

Decoding. We decoded the population response using the Bayes’ least-squares estimator, computed under each model (fully coupled, uncoupled with spike-history terms, and inhomogeneous Poisson) using 6,000 different 18-sample single-pixel stimulus segments (validation data that were not used for fitting). Each stimulus \mathbf{x}_i (an 18-dimensional binary vector, given by the time series of light intensities for a centrally located stimulus pixel) was decoded by first extracting \mathbf{y}_i , the multi-neuronal spike response portion that was causally influenced by this stimulus. For each model, and for every one of the 2^{18} possible binary \mathbf{x}_i , we then computed $p_j = p(\mathbf{y}_i | \mathbf{x}_i)$, the likelihood of the observed population response given that it was generated by stimulus \mathbf{x}_i . By Bayes’ rule, the posterior is $p(\mathbf{x}_i | \mathbf{y}_i) \propto p(\mathbf{y}_i | \mathbf{x}_i) p(\mathbf{x}_i)$, and the prior $p(\mathbf{x}_i)$ here is constant across binary stimuli. Thus, the posterior is proportional to p_j , and the Bayes’ least-squares estimate is given by $\hat{\mathbf{x}}_i = (\sum p_j \mathbf{x}_j) / (\sum p_j)$. We also performed decoding on longer (30-sample) stimulus segments, where exhaustive evaluation of these sums is no longer tractable: in this case we used Gibbs sampling from $p(\mathbf{x}_j | \mathbf{y}_j)$ to approximately evaluate the sum. The results obtained using both methods were similar.

Linear decoding was performed using the optimal linear estimator⁶, with the same training data as for model fitting. Decoding performance was quantified using the log SNR of each technique: $\log \left(\frac{\langle \mathbf{x}_i \mathbf{x}_i^T \rangle}{\langle \mathbf{r}_i \mathbf{r}_i^T \rangle} \right)$, where $\mathbf{r}_i = \hat{\mathbf{x}}_i - \mathbf{x}_i$ denotes the residual error for decoding stimulus vector \mathbf{x}_i , and $\langle \cdot \rangle$ denotes averaging over i followed by matrix determinant. Breakdown by temporal frequency was obtained by computing the Fourier power spectra of the stimuli $\hat{\mathbf{x}}_i(\omega)^2$ and residuals $\hat{\mathbf{r}}_i(\omega)^2$, and computing log SNR according to $\log \left(\frac{\langle \hat{\mathbf{x}}_i(\omega)^2 \rangle}{\langle \hat{\mathbf{r}}_i(\omega)^2 \rangle} \right)$. Integrating this log SNR across frequency, $(1/2) \int \log \text{SNR}(\omega) d\omega$, gives a commonly used estimate of the mutual information between the stimulus and the spike-train response⁶, which is equivalent to the quantity shown in Fig. 4b.

30. Meister, M. & Berry, M. J. The neural code of the retina. *Neuron* 22, 435–450 (1999).