

Multiple neural spike train data analysis: state-of-the-art and future challenges

Emery N Brown, Robert E Kass & Partha P Mitra

Multiple electrodes are now a standard tool in neuroscience research that make it possible to study the simultaneous activity of several neurons in a given brain region or across different regions. The data from multi-electrode studies present important analysis challenges that must be resolved for optimal use of these neurophysiological measurements to answer questions about how the brain works. Here we review statistical methods for the analysis of multiple neural spike-train data and discuss future challenges for methodology research.

Neurophysiologists often administer a stimulus and simultaneously record neural activity from a brain region believed to respond to that stimulus. The stimulus can be physical in nature, such as light used to stimulate retinal or lateral geniculate neurons, or sound used to stimulate neural activity in the auditory cortex. It can also be abstract or cognitive, such as in a working memory task, which elicits neural activity in the hippocampus or pre-frontal cortex. The experimental question can be addressed by characterizing the relation between the stimulus and the individual or ensemble neural responses and/or the relation among the spiking activity of the neurons in the ensemble. In contrast to studying the spiking activity from a single neuron, the recent advent of multiple-electrode recording¹ makes it possible to study the simultaneous spiking activity of many neurons (more than 20). This allows us to understand how groups of neurons act in concert to define the function of a given brain region. Simultaneous recording of multiple neurons offers new promise for investigating fundamental questions, provided the challenging problem of analyzing multiple simultaneously recorded spike trains can be properly addressed.

In probability theory and statistics, a time series of discrete events, such as a spike train, is called a point process². Hence, ensembles of spike trains from simultaneously recorded neurons are multi-dimensional point-process time series. These time series are both dynamic and stochastic. That is, their properties change through time in a manner that can often be characterized by a probability model describing the likelihood of spikes at a given time. These data present

new analysis challenges because most standard signal processing techniques are designed primarily for continuous-valued data and not point processes. Thus, standard methods are of limited use in analyzing multiple neural spike train data. Moreover, because brain regions represent relevant biological signals in the spiking patterns of their constituent neurons, proper analysis of these data requires accurately characterizing the neural interactions.

Spike sorting: identification and classification of spike events

In neurophysiological experiments, individual spikes are not directly recorded. This is because when multiple electrodes are implanted, the extracellular voltage potentials recorded on any electrode represent the simultaneous electrical activity of an unknown number of neurons. From these voltage traces, the spike events or action potentials must be identified, the number of neurons being recorded must be determined, and each spike must be assigned to the neuron that produced it^{3–5}. This three-stage process, termed 'spike sorting' (Fig. 1a,b) is the mandatory first step in all multiple spike train data analyses. The accuracy of the spike sorting critically affects the accuracy of all subsequent analyses.

Many algorithms are used for spike sorting and at present, there is no consensus as to which are best. Different algorithms applied to the same data set can yield different results, illustrating the many complexities of the spike-sorting problem. First, clusters of voltage traces that summarize the spike events often violate the frequently made assumption of stable, Gaussian errors in model-based parametric algorithms. Because neuronal properties and experimental conditions evolve, these clusters change over time. Second, identifying the number of neurons is a challenging problem. One strategy is to assume a number of neurons well in excess of the number believed to be in the data, and then combine clusters that are sufficiently close using a stopping criterion⁵. An alternative Monte Carlo-based strategy has been recently proposed, but has yet to be widely tested⁶. Third, dual intracellular-extracellular recording studies have shown that spike sorting, particularly for large numbers of neurons, has a non-zero error rate because the probability distribution of spike shapes from different neurons share some degree of overlap⁴. Finally, multiple electrodes with different geometries and numbers of electrodes usually require different sorting algorithms.

Cross-correlogram and cross-intensity function

Most current methods for neural spike train data analysis assess only associations between pairs of neurons. As is true for continuous-valued data, techniques to measure the association between neural spike trains can be divided into time-domain and frequency-

Emery N. Brown is in the Neuroscience Statistics Research Laboratory, Department of Anesthesia and Critical Care, Massachusetts General Hospital, and the Division of Health Sciences and Technology, Harvard Medical School, Massachusetts Institute of Technology, Boston, Massachusetts 02114, USA. Robert E. Kass is in the Department of Statistics, Carnegie Mellon University and the Center for the Neural Basis of Cognition, Pittsburgh, Pennsylvania 15208, USA. Partha P. Mitra is at the Cold Spring Harbor Laboratory, Cold Spring Harbor, New York 11724, USA.
e-mail: brown@neurostat.mgh.harvard.edu

Published online 27 April 2004; doi:10.1038/nn1228




Figure 1 Transition from voltage signal recordings to measures of association for three neural spike trains. (a) Voltage trace containing the spike events of three different neurons recorded on the same electrode. Each colored star indicates a different neuron. (b) Application of a spike sorting algorithm that identifies the spike events, determines the number of neurons and assigns each spike event to a particular neuron. The dotted vertical lines show a spike triplet identified by a spike pattern classification method. (c–f) Measures of association between the spike trains from Cell 1 and Cell 2 computed using an unnormalized cross-correlogram (c), a JPSTH (d), a parametric model fit by maximum likelihood (e) and a cross-coherence function (f, solid black line) and confidence bounds (f, thin black line). The horizontal line in c is the upper 95% confidence bound. Correlations above this line are significantly different from zero.

domain methods. The most commonly used time-domain method for measuring association between neurons is the unnormalized cross-correlogram (Fig. 1c)⁷. Given a pair of neural spike trains and a specified bin width, the un-normalized cross-correlogram is the cross-covariance between the two binned spike trains computed at a series of lags. The method assumes that the two spike trains are stationary. That is, it is assumed that the stochastic properties of the neurons do not change in time. In many cases, this ‘stationarity’ assumption can be hard to justify, given that the neural responses are elicited by time-varying stimuli and frequently adapt with time in response to the same stimulus. Non-stationarity has been addressed by performing the covariance analyses in moving windows; however, this requires a substantial amount of data.

A related measure of association between two stationary point processes (spike trains) is the cross-intensity function⁸. This function estimates the spike rate of one neuron at different lags relative to the

spiking activity of a second neuron. Despite being designed expressly to measure association between two point processes, being simple to compute and having associated confidence interval estimates, this method has received only limited use in neural data analysis. Both the cross-correlogram and the cross-intensity function are histogram-based, and provide only measures of paired associations of neural activity.

Joint peri-stimulus time histogram

The joint peri-stimulus time histogram (JPSTH)⁹ (Fig. 1d) is for a pair of neurons a logical extension of the single-neuron PSTH^{9–11}. Whereas the PSTH displays the spike count per unit time t at each time t , the JPSTH is a two-dimensional histogram that displays the joint spike count per unit time at each time u for neuron 1 and time v for neuron 2. The main diagonal of the JPSTH (the ‘PST coincidence histogram’) displays for each time t the observed rate at which both neurons fire simultaneously (to within the accuracy of the binwidth of the histogram). A modification of the JPSTH, termed the normalized JPSTH, is also used¹². The normalized JPSTH subtracts from the joint firing rate the firing rate expected under independence, and then divides by the product of the two standard deviations (of the two neuronal firing rates) to correct for the possibility that two independent neurons with jointly elevated firing rates can appear to be strongly associated. The normalized JPSTH at the time pair (u, v) is the Pearson correlation (computed across trials) of the firing of neuron 1 at time u with that of neuron 2 at time v . Summing the diagonals of the normalized JPSTH produces the normalized cross-correlogram.

Although the normalized JPSTH and the normalized cross-correlogram (Fig. 1) are useful, both have limitations. First, the Pearson correlation is only one of many possible measures of association, and different measures can produce different results, the accuracy of which depends on the underlying mechanism that produces the joint spiking activity¹³. Second, statistical significance testing can be performed in several ways with these methods and again, the results can differ depending on the assumptions and the methods. A new approach to significance testing using recently developed smoothing procedures and a bootstrap significance test can yield greater statistical power¹¹. The bootstrap is a broadly applicable simulation method for estimating uncertainty in a statistical analysis. Third, the normalized JPSTH and cross-correlogram assume that all trials are statistically indistinguishable⁷. If, instead, there is detectable trial-to-trial variation in the neural firing rates, then this variation can appear artifactually as synchrony or time-lagged joint firing^{7,14}. A fourth, crucial consideration is that whereas all spike train analysis is predicated on good spike sorting, the accuracy of spike time information is particularly important when searching for synchrony or time-lagged joint firing. The effects of spike overlap, which are problematic for most spike-sorting algorithms, can produce spurious correlations between pairs of neurons¹⁵.

Spike pattern classification methods

Algorithms to detect precise patterns of spike timing are another method of measuring associations among neural spike trains^{10,16,17}. The appeal of these methods is that they provide a way of evaluating higher-order neural interactions, that is, greater than pairwise, in ensemble spiking activity¹⁸. For example, these methods can be used to assess the statistical significance of spike triplet occurrences separated by precise interspike intervals or the occurrence of similar patterns among two or more neurons (Fig. 1b)¹⁹. Methods for identifying statistically conspicuous spike coincidences have also been developed. Such coincidences have been labeled ‘unitary events’ when they occur more frequently than would be predicted by chance under the null

hypothesis that spike times are independent^{20–22}. Their occurrence is then studied in relation to behavioral events. The delicate statistical issue involved in applying spike pattern classification methods is choosing the complexity or size of the pattern, and formulation of the null hypothesis and test statistic so that the procedure has the correct significance level under reasonably general assumptions. For this reason, some findings from these analyses have at times been criticized as suffering from statistical artifacts²³. An alternative approach to test for synchrony is to build a distribution of spike trains under a null hypothesis by ‘jittering’ the observed spike trains randomly within a small time window. This intuitive idea has recently been formalized and extended to cover several practical data analysis scenarios²⁴.

Likelihood methods

Likelihood methods are central tools for modeling and analysis in statistical research²⁵. Most likelihood methods assume a specific parametric probability model for a process under study (Fig. 1e). The likelihood is the joint probability density of the experimental data arising from this process viewed as a function of the model’s unknown or free parameters. These free parameters may be estimated from the experimental data by formal estimation procedures such as method of moments or maximum likelihood. If the probability model is a good approximation to the process being studied, then use of the likelihood is an optimal way of analyzing the data being generated by the process²⁵. Likelihood methods for point processes have been used to analyze single neural spike train data^{8,26–29}, and in a few instances to model two or more simultaneously recorded neurons^{8,30}. Likelihood methods hold important promise for this and other neuroscience data analysis problems because they provide in a coherent framework a wide range of well-developed statistical methods for data analysis, including assessing model goodness-of-fit, constructing confidence intervals and testing hypotheses^{8,29}. The challenge in using likelihood point process methods to analyze multiple neural spike trains is defining multivariate point process models that accurately represent joint neural spiking activity and devising efficient algorithms for model fitting³⁰.

Frequency-domain methods

Under the assumption of stationarity, as in the case of continuous-valued data, a frequency domain analysis of ensemble neural spiking activity can be conducted by taking the Fourier transform of the spike trains, and using these to compute the spectrum of the individual trains and the cross-spectrum or coherence between each spike train pair^{8,31,32}. The coherence is a simple frequency-dependent correlation measure of association between two processes (Fig. 1f). It has two important advantages over the time domain counterpart: the normalization is not bin-size dependent, and it can be pooled across neuron pairs. It also allows for analysis of point processes, continuous-valued processes, and hybrid point and continuous-valued pairs using the same measure. Error estimates and confidence intervals can be computed for spectra and coherence estimates from theoretical formulae

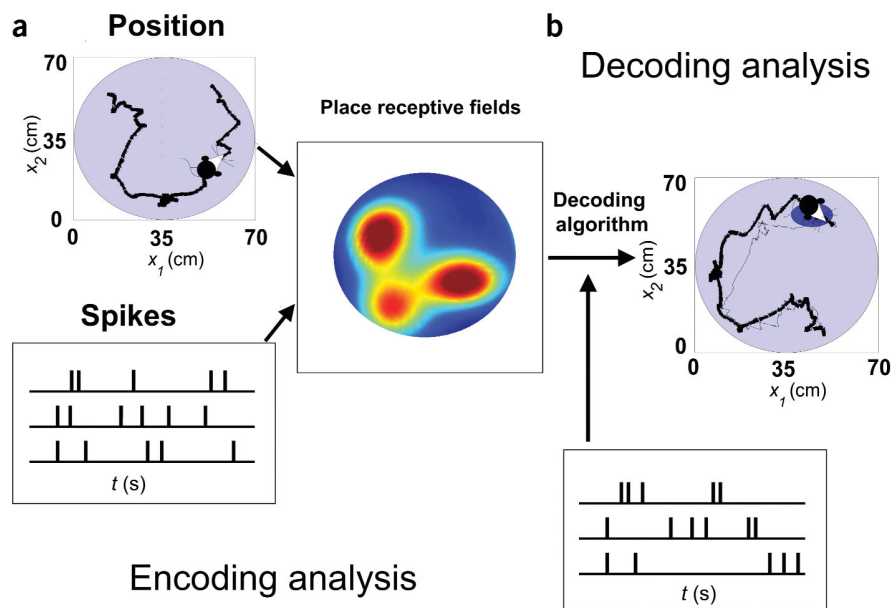


Figure 2 Decoding of position from ensemble rat neural spiking activity^{40,50}. (a) Encoding analysis in which the relation between the biological stimulus (trajectory of the rat in the environment, solid black line in the Position panel) and spiking activity (Spikes panel) is estimated as place receptive fields for three neurons. (b) Decoding analysis in which the estimated place receptive fields are used in a Bayesian decoding algorithm to compute the predicted position (thin black line) of the rat in the environment from new spiking activity of the neural ensemble recorded during the decoding stage. The predicted position is compared with the observed position (thick black line) during the decoding stage. The blue oval defines a 95% confidence region centered at that location.

that are valid when the numbers of spikes in the spike trains are large, or from bootstrap/jackknife procedures (Fig. 1f)³³.

Stimulus-driven non-stationarity is an important feature of neural spike train data, and may be analyzed using moving window estimates of spectra (spectrograms) and coherences (coherograms)³⁴. A key technical yet practical point for use of time-frequency spectral estimates, including moving window and wavelet-based estimates³⁵, is that they must obey the uncertainty principle, which puts a lower bound on the area of the point spread functions of these estimates at all points in the time-frequency plane ($\Delta f \Delta t \geq 1$). Moving window estimates computed in the frequency domain are often less biased than the corresponding time-domain estimates. Thus, even time domain functions, such as the cross-correlogram and the PSTH, may be optimally estimated by inverse Fourier-transforming the corresponding frequency-domain quantities. One principled approach to estimating the frequency-domain quantities is by using multitaper techniques³⁶. These methods have also proved useful in estimating coherence between spike trains and local field potentials³⁷ and are well-suited for error analyses using bootstrap/jackknife procedures.

Neural spike train decoding

Decoding algorithms are the mathematical techniques used in neuroscience to study how spike train firing patterns from a single neuron^{38,39} or an ensemble of neurons⁴⁰ represent external stimuli and biological signals. The decoding analysis proceeds typically in two stages: the encoding stage (Fig. 2a) and the decoding stage (Fig. 2b). In the encoding stage, neural spiking activity is characterized as a function of the biological signal. In the decoding stage, the relation is inverted, and the signal is estimated from the spiking activity of the neurons. Developed initially to study how movements are represented

by neurons in the motor cortex, the population vector is one of the earliest decoding techniques^{41,42}. To decode the neural spiking activity at any time t , one computes the population vector (the normalized dot product of the observed spiking activity in a time window with the firing functions for each neuron) for different values of the signal. The value of the signal for which the dot product is the largest is taken to be the decoded estimate of the signal.

Reverse correlation is a linear regression–based decoding algorithm that has been used to study how groups of neurons represent information in the visual and motor systems and to control neural prosthetic devices^{38,39,43–46}. The appeal of this widely used approach is that binning the spike trains to create continuous-valued regressors avoids use of an explicit encoding model and makes it possible to use standard linear regression theory to fit the model and assess the accuracy of the decoding. Moreover, computation of the regression coefficients implicitly takes account of pairwise correlations in the neural activity. The linear construction of the population vector makes it a special case of the reverse correlation methods.

Decoding algorithms that are based on Bayes' theorem, the elementary probability rule for computing the probability of one event given another event, offer a general approach to estimating the representation of a biological signal in ensemble spiking activity^{40,47–51}. They have been used successfully to study how neural ensembles in the hippocampus represent an animal's position in an environment^{40,47,50} (Fig. 2) and to characterize how motor commands are represented by ensembles of neurons in primary motor cortex^{49,51}. The encoding stage can use likelihood methods to compute the probability of the spiking activity given the signal, and the decoding stage computes the probability of the signal given the spiking activity. The appeal of the Bayesian approach is that it uses probability models to represent different sources of information in the problem, and it formulates decoding in the theoretical framework of other filtering and smoothing methods in statistics and signal processing. When the proposed model is a reasonable approximation to the data, the Bayesian approach, like the likelihood methods, has many optimality properties, including efficiency, which, in the decoding problem, means that its signal estimates have the smallest possible uncertainty²⁵. An important conceptual difference between the Bayesian and reverse correlation decoding methods is that under the standard assumptions of regression theory, the neural firing rates used in a reverse correlation analysis are assumed to be non-random, known constants. In contrast, the Bayesian approach models the spike trains as a stochastic point process and the biological signal as a stochastic process based on its known properties.

Information theory

Information theory measures are used widely in analyses of neural spike train data^{39,52–54}. These include the entropy to quantify spike train variability, and mutual information to measure the association between two processes, such as between two spike trains or between a spike train and a stimulus. These measures have been applied extensively to study how much information a single spike train conveys about a biological signal by using histogram-based methods to estimate empirically the relevant probability densities. Use of the information measures is grounded in thinking about those parts of the nervous system, such as visual pathways, that may be modeled as communication channels³⁹ with a rationale that analyses may be conducted free of assumptions about detailed system properties⁵⁵. There are limitations to this approach. For any neural system, the optimal word length (histogram binwidth) is an unknown that must be estimated taking account that the data requirements for histogram esti-

mation increase exponentially with the word length. The data requirements are far greater for extending this approach to estimating mutual information between multiple spike trains and a biological signal⁵⁴. Information theory methods summarize complex functional relationships between the spike train and the signal as single numbers. Moreover, whether sensory pathways can be treated using information theory as in conventional communications analyses has recently been questioned⁵⁶.

One approach to extending the use of information theory to analyze multiple spike trains may arise from developing probability models of joint spiking activity and likelihood methods to estimate these models. An advantage of modeling explicitly the joint probability density between the ensemble spiking activity and the biological signal is that the mutual information and any other relevant functions of this probability density can be computed directly once this probability density has been estimated. Parametric models may then offer insight into how to construct their more flexible, model-free counterparts.

Future challenges for multiple spike train data analysis

Simultaneous recording of multiple spike trains from several neurons offers a window into how neurons work in concert to generate specific brain functions. Without substantial methodology research in the future, our ability to understand this function will be significantly hampered because current methods fall short of what is ultimately required for the analysis of multiple spike train data. With the exception of spike pattern classification methods, decoding algorithms, partial coherence estimation⁸ and certain graphical methods⁵⁷, current techniques for spike train analysis are designed to analyze—at most—pairs of neurons. Therefore, the future challenge is to design methods that truly allow neuroscientists to perform multivariate analyses of multiple spike train data. This development must be done taking explicit account of the questions being studied and the experimental protocols being used.

Because the accuracy of the spike sorting significantly affects the accuracy of the experimental data, development of the best possible spike-sorting algorithms must be an important goal. The complexity of the spike-sorting problem increases with number of electrodes in the recording systems. There should be systemic study of spike-sorting algorithms taking account of different electrode numbers and configurations, recording conditions and brain regions. A harder, yet no less important, challenge is to devise accurate, real-time spike-sorting algorithms to enable multiple spike trains to be inputs to neural prosthetic devices or brain-machine interfaces^{42,45,46}. Real-time spike sorting could also lead to real-time data analysis, and possibly to real-time changes in experimental protocols.

Graphical methods⁵⁷ for multivariate point process data are important for screening data for errors and inconsistencies prior to analysis, postulating preliminary models and formulating meaningful displays to report findings.

Multiple spike trains are multivariate point processes, yet research in statistics and signal processing on multivariate point process models has not been nearly as extensive as research on models of multivariate continuous-valued processes. Therefore, developing multivariate point process models should be a primary focus of methodology research for multiple neural spike trains. Because there is a canonical representation of univariate and multivariate point processes in terms of the conditional intensity function^{2,29}, developing strategies to construct parametric models of conditional intensity functions and likelihood-based estimation methods seems a good way to proceed^{2,8,27–30}. Other avenues of investigation could include lattice or spin models from statistical physics⁵⁸ and multivariate binary data models from

statistics⁵⁹. Whatever the approach, the objective must be to develop tractable methods for estimating high-dimensional interactions among groups of neurons from their spike train recordings. Furthermore, because plasticity in neural dynamics makes non-stationarity in neural data a rule rather than an exception, developing explicit adaptive estimation algorithms to track these dynamics for multivariate point processes is another important research problem⁶⁰.

Dynamical systems and neural network models have long been central in providing quantitative characterizations of neural processes^{61,62}. Research on data analysis methodology should be conducted in concert with this modeling research⁶³. Multiple-electrode recordings combined with statistical methods explicitly designed to analyze multiple spike train data will offer a better opportunity to explicitly link experimentation and computational modeling by using the information from experiments to quantify better predictions from more complex models, refine model formulation and, as a consequence, design better experiments. Similarly, the computational models can suggest formulations of the statistical methods that may enhance their success at extracting salient features in experimental data.

Although the objective of most current neurophysiological experiments is to relate relevant biological stimuli to ensemble spiking activity, experiments that record simultaneous multimodality data such as neurophysiological, functional imaging and behavioral data are becoming more common⁶⁴. Developing appropriate statistical methods to analyze simultaneous multimodality recordings will require innovative approaches to integrate information properly across the different temporal and spatial scales of various data sources.

There are many benefits of developing multivariate methods for multiple spike train data analysis. First, methods specifically tailored to analyze multiple spike train data will allow neuroscientists to make precise statements of how reliably findings from a given experiment can be stated in terms of standard statistical summaries. Practically speaking, this means that even for this complex, high-dimensional modeling problem, the analysis reports standard errors for firing rates and time constant parameters, provides confidence intervals for measures of neural interactions and associations, and gives quantitative assessments of how well a given model describes the experimental data²⁹. Second, more accurate quantitative summaries will allow neuroscientists to make more reliable statements about how strongly experimental findings support hypotheses or proposed mechanisms. For example, would an analysis measuring time-varying interactions among three or more neurons rather than pairwise correlations offer new insight into the mechanism of persistent activity seen in the oculomotor system⁶⁵?

Third, more accurate multivariate quantitative summaries will make it easier to relate ensemble neural dynamics (within and between specific brain regions) to behavior and to relevant biological stimuli. As an illustration, applying these methods to the study of simultaneously recorded neural activity in the parietal and primary motor cortices could help reveal how these two brain regions communicate during formulation and execution of motor commands. Fourth, as the number of neurons whose interactions can be accurately measured increases, neuroscientists will be able to increase the complexity of the experiments they design, and as a consequence, the questions they investigate. Fifth, more reliable data analyses will provide more refined quantitative constraints and perhaps parameter values for dynamical models of neural systems. Finally, improved multiple spike train data analysis methods, particularly spike-sorting and decoding algorithms, will have immediate, significant implications for improving the design and implementation of neural prosthetic devices and brain-computer interfaces^{42,45,46}.

Multiple spike train recordings are an important component of the data explosion that is currently occurring in neuroscience. Therefore, devising systematic research programs for neuroscience data analysis akin to those currently being undertaken in genomics and bioinformatics is a must. Several initiatives to support such research have already been proposed by the US National Institutes of Health (<http://grants1.nih.gov/grants/guide/pa-files/PA-04-006.html>) and National Science Foundation (www.nsf.gov/bio/progdes/biocrcn.htm). Specific initiatives to encourage quantitative scientists (statisticians, physicists, engineers, computer scientists and mathematicians) to undertake data analysis research in neuroscience should be part of these current and future programs. Neuroscience training for statisticians and incentives to involve them more directly in neuroscience data analysis research should be a priority. Finally, courses on the analysis of neuroscientific data (www.mbl.edu/education/courses/special_topics/neufo.html) should be part of the curriculum in neuroscience programs, as are courses on computational modeling. This will ensure that instruction in the most contemporary data analysis principles and methods are an integral part of undergraduate, graduate and postdoctoral training in neuroscience, and in the disciplines that support computational research in this field.

ACKNOWLEDGMENTS

Support for this work was provided in part by NIH grants MH66410 to P.M. and E.N.B., MH62528 to P.M., MH64537 to R.E.K., and MH59733, MH61637 and DA015664 to E.N.B. We thank S. Grün for comments on an earlier draft of this manuscript, G. Gerstein for permission to use Fig. 1d and R. Barbieri for help preparing the figures.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>

- Wilson, M.A. & McNaughton, B.L. Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058 (1993).
- Daley, D. & Vere-Jones, D. *An Introduction to the Theory of Point Process*, 2nd ed. (Springer-Verlag, New York, 2003).
- Lewicki, M.S. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network Comput. Neural Syst.* **9**, R53–R78 (1998).
- Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H. & Buzsáki, G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.* **84**, 401–414 (2000).
- Fee, M.S., Mitra, P.P. & Kleinfeld, D. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J. Neurosci. Methods* **69**, 175–188 (1996).
- Nguyen, D., Frank, L.M. & Brown, E.N. An application of reversible-jump MCMC to spike classification of multiunit extracellular recordings. *Network Comput. Neural Syst.* **14**, 61–82 (2003).
- Brody, C.D. Correlations without synchrony. *Neural Comput.* **11**, 1537–1551 (1999).
- Brillinger, D.R. Nerve cell spike train data analysis: a progression of techniques. *J. Amer. Stat. Assoc.* **87**, 260–271 (1992).
- Gerstein, G.L. & Perkel, D.H. Simultaneously recorded trains of action potentials: analysis and functional interpretation. *Science* **164**, 828–830 (1969).
- Abeles, M. Quantification, smoothing, and confidence limits for single-unit histograms. *J. Neurosci. Methods* **5**, 317–325 (1982).
- Kass, R.E., Ventura, V. & Cai, C. Statistical smoothing of neuronal data. *Network Comput. Neural Syst.* **14**, 5–15 (2003).
- Aertsen, A., Gerstein, G.L., Habib, M.K. & Palm, G. Dynamics of neural firing correlation: modulation of "effective connectivity." *J. Neurophysiol.* **61**, 900–917 (1989).
- Ito, H. & Tsujii, S. Model dependence in quantification of spike interdependency by joint peri-stimulus time histogram. *Neural Comput.* **12**, 195–217 (2000).
- Baker, S.N. & Gerstein, G.L. Determination of response latency and its application to normalization of cross-correlation measures. *Neural Comput.* **13**, 1351–1377 (2001).
- Bar-Gad, I., Ritov, Y., Vaadia, E. & Bergmann, H. Failure in identification of overlapping spikes from multiple neuron activity causes artificial correlations. *J. Neurosci. Methods* **107**, 1–13 (2001).
- Softky, W.R. & Koch, C. The highly irregular firing of cortical cells is consistent with temporal integration of random EPSPs. *J. Neurosci.* **13**, 334–350 (1993).
- Shadlen, M.T. & Newsome, W.N. The variable discharge of cortical neurons: impli-

- cations for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
18. Martignou, L. *et al.* Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput.* **12**, 2621–2653 (2000).
 19. Abeles, M. & Gerstein, G.L. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.* **60**, 909–924 (1988).
 20. Grün, S., Diesmann, M. & Aertsen, A. Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput.* **14**, 81–119 (2002).
 21. Güttig, R., Aertsen, A. & Rotter, S. Statistical significance of coincident spikes: count-based versus rate-based statistics. *Neural Comput.* **14**, 121–153 (2002).
 22. Pipa, G. & Grün, S. Non-parametric significance estimation of joint-spike events by shuffling and resampling. *Neurocomputing* **52–54**, 31–37 (2003).
 23. Oram, M.W., Wiener, M.C., Lestienne, R. & Richmond, B.J. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J. Neurophysiol.* **81**, 3021–3033 (1999).
 24. Amarasingham, A. *Temporal Structure in Nervous System Activity* dissertation, Brown Univ. (2002).
 25. Pawitan, Y. *In All Likelihood: Statistical Modelling and Inference Using Likelihood* (Oxford Univ. Press, New York, 2001).
 26. Brillinger, D.R. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol. Cybern.* **59**, 189–200 (1988).
 27. Barbieri, R., Quirk, M.C., Frank, L.M., Wilson, M.A. & Brown, E.N. Construction and analysis of non-Poisson stimulus response models of neural spike train activity. *J. Neurosci. Methods* **105**, 25–37 (2001).
 28. Kass, R.E. & Ventura, V. A spike train probability model. *Neural Comput.* **13**, 1713–1720 (2001).
 29. Brown, E.N., Barbieri, R., Eden, U.T. & Frank, L.M. Likelihood methods for neural data analysis. in *Computational Neuroscience: A Comprehensive Approach* (ed. Feng, J.) 253–286 (CRC Press, Boca Raton, 2003).
 30. Chornoboy, E.S., Schramm, L.P. & Karr, A.F. Maximum likelihood identification of neural point process systems. *Biol. Cybern.* **59**, 265–275 (1988).
 31. Brillinger, D.R. Comparative aspects of the study of ordinary time series and of point processes. in *Developments in Statistics* Vol. 1, 33–129 (Academic Press, Orlando, 1978).
 32. Jarvis, M.R. & Mitra, P.P. Sampling properties of the spectrum and coherency in sequences of action potentials. *Neural Comput.* **13**, 717–749 (2001).
 33. Thomson, D.J. & Chave, A.D. Jackknifed error estimates for spectra, coherences, and transfer functions. in *Advances in Spectrum Analysis and Array Processing* (ed., Haykin, S.) 58–113 (Prentice Hall, Englewood Cliffs, NJ, 1991).
 34. Brillinger, D.R. *Time Series* (Holt, Rinehart and Winston, New York, 1981).
 35. Percival, D.B. & Walden, A.T. *Wavelet Methods for Time Series Analysis* (Cambridge Univ. Press, Cambridge, UK, 2000).
 36. Percival, D.B. & Walden, A.T. *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques* (Cambridge Univ. Press, Cambridge, UK, 2002).
 37. Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P. & Andersen, R.A. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.* **5**, 805–811 (2002).
 38. Bialek, W., Rieke, F., de Ruyter van Steveninck, R.R. & Warland, D. Reading a neural code. *Science* **252**, 1854–1857 (1991).
 39. Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, 1997).
 40. Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C. & Wilson, M.A. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J. Neurosci.* **18**, 7411–7425 (1998).
 41. Georgopoulos, A.P., Kettner, R.E. & Schwartz, A.B. Neuronal population coding of movement direction. *Science* **233**, 1416–1419 (1986).
 42. Taylor, D.M., Tillery, S.I. & Schwartz, A.B. Direct cortical control of 3D neuroprosthetic devices. *Science* **296**, 1829–1832 (2002).
 43. Warland, D.K., Reinagel, P. & Meister, M. Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* **78**, 2336–2350 (1997).
 44. Stanley, G.B., Li, F.F. & Dan, Y. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *J. Neurosci.* **19**, 8036–8042 (1999).
 45. Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* **408**, 361–365 (2000).
 46. Serruya, M.D., Hatsopoulos, N.G., Paninski, L., Fellows, M.R. & Donoghue, J.P. Instant neural control of a movement signal. *Nature* **416**, 141–142 (2002).
 47. Zhang, K., Ginzburg, I., McNaughton, B.L. & Sejnowski, T.J. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* **79**, 1017–1044 (1998).
 48. Wiener, M.C. & Richmond, B.J. Decoding spike trains instant by instant using order statistics and the mixture-of-Poissons model. *J. Neurosci.* **23**, 2394–2406 (2003).
 49. Gao, Y., Black, M.J., Bienenstock, E., Wu, W. & Donoghue, J.P. A quantitative comparison of linear and non-linear models of motor cortical activity for the encoding and decoding of arm motions. *First Intl. IEEE/EMBS Conf. on Neural Eng.* 189–192 (2003).
 50. Barbieri, R. *et al.* Dynamic analyses of information encoding by neural ensembles. *Neural Comput.* **16**, 277–307 (2004).
 51. Brockwell, A.E., Rojas, A.L. & Kass, R.E. Recursive Bayesian decoding of motor cortical signals by particle filtering. *J. Neurophysiol.* **91**, 1899–1907 (2004).
 52. Borst, A. & Theunissen, F.E. Information theory and neural coding. *Nat. Neurosci.* **2**, 947–957 (1999).
 53. Reich, D.S., Melcher, F. & Victor, J.D. Independent and redundant information in nearby cortical neurons. *Science* **294**, 2566–2568 (2001).
 54. Nirenberg, S., Carcieri, S.M., Jacobs, A.L. & Latham, P.E. Retinal ganglion cells act largely as independent encoders. *Nature* **411**, 698–701 (2001).
 55. Strong, S.P., Koberle, R., de Ruyter van Steveninck, R.R. & Bialek, W. Entropy and information in neural spike trains. *Phys. Rev. Lett.* **80**, 197–200 (1998).
 56. Berger, T. Living information theory. *IEEE Information Theory Society Newsletter* **53**, 1–20 (2003).
 57. Stuart, L., Walter, M. & Borisjuk, R. Visualization of synchronous firing in multi-dimensional spike trains. *BioSystems* **67**, 265–279 (2002).
 58. Abdollahi, L.M., La Rota, C., Beguin, M. & François, O. Parameter estimation in a model for multidimensional recording of neuronal data: a Gibbsian approximation approach. *Biol. Cybern.* **89**, 170–178 (2003).
 59. Quaquish, B.F. A family of multivariate binary distributions for simulating correlated binary variables with specified marginal means. *Biometrika* **90**, 455–464 (2003).
 60. Brown, E.N., Nguyen, D.P., Frank, L.M., Wilson, M.A. & Solo, V. An analysis of neural receptive field plasticity by point process adaptive filtering. *Proc. Natl. Acad. Sci. USA* **98**, 12261–12266 (2001).
 61. Dayan, P. & Abbott, L.F. *Theoretical Neuroscience* (MIT Press, Cambridge, 2001).
 62. Shikri, O., Hansel, D. & Sompolinsky, H. Rate models for conductance-based cortical neuronal networks. *Neural Comput.* **5**, 1809–1841 (2003).
 63. Victor, J.D. & Brown, E.N. Information and statistical structure in spike trains. *Network Comput. Neural Syst.* **14**, 1–4 (2003).
 64. Logothetis, N.K., Pauls, J., Augath, M., Trinath, T. & Oeltermann, A. Neurophysiological investigation of the basis of the fMRI signal. *Nature* **412**, 150–157 (2001).
 65. Aksay, E., Baker, R., Seung, H.S. & Tank, D.W. Correlated discharge among pairs within the oculomotor horizontal velocity-to-position integrator. *J. Neurosci.* **23**, 10852–10858 (2003).