J Neurophysiol 93: 1074–1089, 2005.
First published September 8, 2004; doi:10.1152/jn.00697.2004.

A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects

Wilson Truccolo,[1] Uri T. Eden, [2,3] Matthew R. Fellows,[1] John P. Donoghue,[1] and Emery N. Brown[2,3]
[1]*Neuroscience Department, Brown University, Providence, Rhode Island;* [2]*Neuroscience Statistics Research Laboratory, Department of Anesthesia and Critical Care, Massachusetts General Hospital, Boston; and* [3]*Division of Health Sciences and Technology, Harvard Medical School/Massachusetts Institute of Technology, Cambridge, Massachusetts*

Submitted 8 July 2004; accepted in final form 23 August 2004

Truccolo, Wilson, Uri T. Eden, Matthew R. Fellows, John P. Donoghue, and Emery N. Brown. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 93: 1074–1089, 2005. First published September 8, 2004; doi:10.1152/jn.00697.2004. Multiple factors simultaneously affect the spiking activity of individual neurons. Determining the effects and relative importance of these factors is a challenging problem in neurophysiology. We propose a statistical framework based on the point process likelihood function to relate a neuron's spiking probability to three typical covariates: the neuron's own spiking history, concurrent ensemble activity, and extrinsic covariates such as stimuli or behavior. The framework uses parametric models of the conditional intensity function to define a neuron's spiking probability in terms of the covariates. The discrete time likelihood function for point processes is used to carry out model fitting and model analysis. We show that, by modeling the logarithm of the conditional intensity function as a linear combination of functions of the covariates, the discrete time point process likelihood function is readily analyzed in the generalized linear model (GLM) framework. We illustrate our approach for both GLM and non-GLM likelihood functions using simulated data and multivariate single-unit activity data simultaneously recorded from the motor cortex of a monkey performing a visuomotor pursuit-tracking task. The point process framework provides a flexible, computationally efficient approach for maximum likelihood estimation, goodness-of-fit assessment, residual analysis, model selection, and neural decoding. The framework thus allows for the formulation and analysis of point process models of neural spiking activity that readily capture the simultaneous effects of multiple covariates and enables the assessment of their relative importance.

INTRODUCTION

Understanding what makes a neuron spike is a challenging problem, whose solution is critical for deciphering the nature of computation in single cells and neural ensembles. Multiple factors simultaneously affect spiking activity of single neurons and thus assessing the effects and relative importance of each factor creates the challenge. Neural activity is often studied in relation to 3 types of covariates. First, spiking activity is associated with extrinsic covariates such as sensory stimuli and behavior. For example, the spiking activity of neurons in the rat hippocampus is associated with the animal's position in its environment, the theta rhythm, theta phase precession, and the animal's running velocity (Frank et al. 2002; Mehta et al. 1997, 2000; O'Keefe and Dostrovsky 1971; O'Keefe and Recce 1993). Retinal neurons respond to light intensity and light contrast, and V1 neurons are influenced by the spatiotemporal structure outside their classic receptive fields (Knierim and Vanessen 1992; Sillito et al. 1995; Vinje and Gallant 2000). The spiking activity of neurons in the arm region of the primary motor cortex (MI) is strongly associated with several covariates of motor behavior such as hand position, velocity, acceleration, and generated forces (Ashe and Georgopoulos 1994; Fu et al. 1995; Scott 2003). Second, the current spiking activity of a neuron is also related to its past activity, reflecting biophysical properties such as refractoriness and rebound excitation or inhibition (Hille 2001; Keat et al. 2001; Wilson 1999).

Third, current capabilities to record the simultaneous activity of multiple single neurons (Csicsvari et al. 2003; Donoghue 2002; Nicolelis et al. 2003; Wilson and McNaughton 1993) make it possible to study the extent to which spiking activity in a given neuron is related to concurrent ensemble spiking activity (Grammont and Riehle 1999, 2003; Hatsopoulos et al. 1998, 2003; Jackson et al. 2003; Maynard et al. 1999; Sanes and Truccolo 2003). Therefore, a statistical modeling framework that allows the analysis of the simultaneous effects of extrinsic covariates, spiking history, and concurrent neural ensemble activity would be highly desirable.

Current studies investigating the relation between spiking activity and these 3 covariate types have used primarily linear (reverse correlation) or nonlinear regression methods (e.g., Ashe and Georgopoulos 1994; Fu et al. 1995; Luczak et al. 2004). Although these methods have played an important role in characterizing the spiking properties in many neural systems, 3 important shortcomings have not been fully addressed. First, neural spike trains form a sequence of discrete events or point process time series (Brillinger 1988). Standard linear or nonlinear regression methods are designed for the analysis of continuous-valued data and not point process observations. To model spike trains with conventional regression methods the data are frequently smoothed or binned, a preprocessing step that can alter their stochastic structure and, as a consequence, the inferences made from their analysis. Second, although it is accepted that extrinsic covariates, spiking history, and neural ensemble activity affect neural spiking, current approaches make separate assessments of these effects, thereby making it difficult to establish their relative importance. Third, model goodness-of-fit assessments as well as the analysis of neural

Address for reprint requests and other correspondence: W. Truccolo, Neuroscience Department, Brown University, 190 Thayer St., Providence, RI (E-mail Wilson_Truccolo@Brown.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

ensemble representation based on decoding should be carried out using methods appropriate for the point process nature of neural spike trains.

To address these issues, we present a point process likelihood framework to analyze the simultaneous effects and relative importance of spiking history, neural ensemble, and extrinsic covariates. We show that this likelihood analysis can be efficiently conducted by representing the logarithm of the point process conditional intensity function in terms of linear combinations of general functions of the covariates and then using the discrete time point process likelihood function to fit the model to spike train data in the generalized linear model (GLM) framework. Because the discrete time point process likelihood function is general, we also show how it may be used to relate covariates to neural spike trains in a non-GLM model. We illustrate the methods in the analysis of a simulated data example and an example in which multiple single neurons are recorded from MI in a monkey performing a visuomotor pursuit-tracking task.

METHODS

In this section we present the statistical theory underlying our approach. First, we define the conditional intensity function for a point process. Second, we present a discrete time approximation to the continuous time point process likelihood function, expressed in terms of the conditional intensity function. Third, we show that when the logarithm of the conditional intensity is a linear combination of functions of the covariates, the discrete time point process likelihood function is equivalent to the likelihood of a GLM under a Poisson distribution and log link function. Alternatively, if the point process is represented as a conditionally independent Bernoulli process and the probability of the events is modeled by a logistic function, then the likelihood function is equivalent to the likelihood of a GLM under a Bernoulli distribution and a logistic link function. Fourth, we present several forms of conditional intensity models for representing spiking history, neural ensemble, and extrinsic covariate effects. Finally, we define our approach to maximum likelihood estimation, goodness-of-fit assessment, model comparison, residuals analysis, and decoding from point process observations by combining the GLM framework with analysis methods for point processes.

A *point process* is a set of discrete events that occur in continuous time. For a neural spike train this would be the set of individual spike times. Given an observation interval $(0, T]$, a sequence of J spike times $0 < u_1 < \cdots < u_j < \cdots < u_J \leq T$ constitutes a point process. Let $N(t)$ denote the number of spikes counted in the time interval $(0, t]$ for $t \in (0, T]$. We define a single realization of the point process during the time interval $(0, t]$ as $N_{0:t} = \{0 < u_1 < u_2 < \cdots < u_j \leq t \cap N(t) = j\}$ for $j \leq J$.

Conditional intensity function

A stochastic neural point process can be completely characterized by its conditional intensity function $\lambda(t \mid H(t))$ (Daley and Vere-Jones 2003), defined as

$$\lambda(t \mid H(t)) = \lim_{\Delta \to 0} \frac{P[N(t+\Delta) - N(t) = 1 \mid H(t)]}{\Delta} \quad (1)$$

where $P[\cdot \mid \cdot]$ is a conditional probability and $H(t)$ includes the neuron's spiking history up to time t and other relevant covariates. The conditional intensity is a strictly positive function that gives a history-dependent generalization of the rate function of a Poisson process. From *Eq. 1* we have that, for small Δ, $\lambda(t \mid H(t))\Delta$ gives approximately the neuron's spiking probability in the time interval $(t, t + \Delta]$. Because defining the conditional intensity function completely defines the point process, to model the neural spike train in terms of a point process it suffices to define its conditional intensity function. We use parametric models to express the conditional intensity as a function of covariates of interest, therefore relating the neuron's spiking probability to the covariates. We use $\lambda(t \mid \theta, H(t))$ to denote the parametric representation of the conditional intensity function in *Eq. 1*, where θ denotes an unknown parameter to be estimated. The dimension of θ depends on the form of the model used to define the conditional intensity function.

A discrete time representation of the point process will facilitate the definition of the point process likelihood function and the construction of our estimation algorithms. To obtain this representation, we choose a large integer K and partition the observation interval $(0, T]$ into K subintervals $(t_{k-1}, t_k]_{k=1}^{K}$ each of length $\Delta = TK^{-1}$. We choose large K so that there is at most one spike per subinterval. The discrete time versions of the continuous time variables are now denoted as $N_k = N(t_k)$, $N_{1:k} = N_{0:t_k}$, and $H_k = H(t_k)$. Because we chose large K, the differences $\Delta N_k = N_k - N_{k-1}$ define the spike train as a binary time series of zeros and ones. In discrete time, the parametric form of the conditional intensity function becomes $\lambda(t_k \mid \theta, H_k)$.

Point process likelihood and GLM framework

Because of its several optimality properties, we choose a likelihood approach (Pawitan 2001) for fitting and analyzing the parametric models of the conditional intensity function. As in all likelihood analyses, the likelihood function for a continuous time point process is formulated by deriving the joint probability density of the spike train, which is the joint probability density of the J spike times $0 < u_1 < u_2 < \cdots < u_J \leq T$ in $(0, T]$. For any point process model satisfying *Eq. 1*, this probability density can be expressed in terms of the conditional intensity function (Daley and Vere-Jones 2003). Similarly, in the discrete time representation, this joint probability density can be expressed in terms of the joint probability mass function of the discretized spike train (see APPENDIX *Eqs. A1* and *A2*) and is expressed here as a product of conditionally independent Bernoulli events (Andersen et al. 1992; Berman and Turner 1992; Brillinger 1988; Brown et al. 2003)

$$P(N_{1:K} \mid \theta) = \prod_{k=1}^{K} [\lambda(t_k \mid \theta, H_k)\Delta]^{\Delta N_k}[1 - \lambda(t_k \mid \theta, H_k)\Delta]^{1-\Delta N_k} + o(\Delta^J) \quad (2)$$

where the term $o(\Delta^J)$ relates to the probability of observing a spike train with 2 or more spikes in any subinterval $(t_{k-1}, t_k]$. From *Eqs. A3–A5* in the APPENDIX, it follows that *Eq. 2* can be reexpressed as

$$P(N_{1:K} \mid \theta) = \exp\left\{\sum_{k=1}^{K} \log [\lambda(t_k \mid \theta, H_k)\Delta]\Delta N_k - \sum_{k=1}^{K} \lambda(t_k \mid \theta, H_k)\Delta\right\} + o(\Delta^J) \quad (3)$$

If we view *Eq. 3* as a function of θ, given the spike train observations $N_{1:K}$, then this probability mass function defines our discrete time point process likelihood function and we denote it as $L(\theta \mid H_K) = P(N_{1:K} \mid \theta)$. From *Eq. A6* in the APPENDIX, it can be seen that *Eq. 3* is a discrete time approximation to the joint probability density of a continuous time point process.

To develop a computationally tractable and efficient approach to estimating θ we note that for any subinterval $(t_{k-1}, t_k]$, the conditional intensity function is approximately constant so that, by *Eq. 3*, $P(\Delta N_k) = \exp\{\log [\lambda(t_k \mid \theta, H_k)]\Delta N_k - \lambda(t_k \mid \theta, H_k)\Delta\}$ is given by the Poisson probability mass function. Because Δ is small, this is equivalent to the Bernoulli probability $P(\Delta N_k) = [\lambda(t_k \mid \theta, H_k)\Delta]^{\Delta N_k}[1 - \lambda(t_k \mid \theta, H_k)\Delta]^{1-\Delta N_k}$ in *Eq. 2*. If we now express the logarithm of the conditional intensity function as a linear combination of general functions of the covariates

$$\log \lambda(t_k \mid \theta, H_k) = \sum_{i=1}^{q} \theta_i g_i[v_i(t_k + \tau)] \quad (4)$$

where g_i is a general function of a covariate $v_i(t_k)$ at different time lags τ, and q is the dimension of the estimated parameter θ, then *Eq. 3* has the same form as the likelihood function for a GLM under a Poisson probability model and a log link function (see APPENDIX, *Eqs. A7–A8*). Thus, maximum likelihood estimation of model parameters and likelihood analyses can be carried out using the Poisson–GLM framework. Alternatively, if we extend the results in Brillinger (1988), we obtain

$$\log \{[1 - \lambda(t_k \mid \theta, H_k)\Delta]^{-1}[\lambda(t_k \mid \theta, H_k)\Delta]\} = \sum_{i=1}^{q} \theta_i g_i[v_i(t_k + \tau)] \quad (5)$$

then *Eq. 2* has the same form as the likelihood function for a GLM under a Bernoulli probability distribution and a logistic link function (*Eqs. A9* and *A10*). Thus, maximum likelihood estimation of model parameters and likelihood analyses can also be carried out using the Bernoulli–GLM framework (see also Kass and Ventura 2001). In other words, for Δ sufficiently small (i.e., at most one spike per time subinterval), likelihood analyses performed with either the Bernoulli or the Poisson model are equivalent. However, because we are interested in modeling the conditional intensity function directly, instead of the probability of events in our discrete time likelihoods, we used the Poisson–GLM framework in our analyses.

Therefore, we can take advantage of the computational efficiency and robustness of the GLM framework together with all of the analysis tools from the point process theory: goodness-of-fit based on the time rescaling theorem, residual analysis, model selection, and stochastic decoding based on point process observations. We refer to this combined framework as the *point process–GLM framework.* This framework covers a very large class of models because *Eq. 4* allows for general functions of the covariates and of interaction terms consisting of combinations of the covariates. An application of GLM analysis to spike train data, without the support of the derived relations between the point process and GLM likelihood functions, would remain purely heuristic in nature.

Finally, *Eqs. 2* and *3* are generally applicable discrete time approximations for the point process likelihood function. Thus, when a parametric model of the conditional intensity function cannot be expressed in terms of either *Eq. 4* or *Eq. 5*, the GLM framework may be replaced with standard algorithms for computing maximum likelihood estimates (Pawitan 2001).

Models for the conditional intensity function

We formulate specific models for the conditional intensity function that incorporate the effects of spiking history, ensemble, and extrinsic covariates. For the exposition in the remainder of this section, we extend our notation to include the neural ensemble activity. Consider an observation time interval $t \in (0, T]$ with corresponding sequences of J^c spike times $0 < u_1^c < \cdots < u_j^c < \cdots < u_{J^c}^c \leq T$, for $c = 1, \ldots, C$ recorded neurons. Let $N_{1:K}^{1:C} = \cup_{c=1}^{C} N_{1:K}^c$ denote the sample path for the entire ensemble.

CONDITIONAL INTENSITY MODELS IN THE POINT PROCESS–GLM FRAMEWORK. The general form for the conditional intensity function we use to model a single cell's spiking activity is

$$\lambda(t_k \mid N_{1:k}^{1:C}, \mathbf{x}_{k+\tau}, \theta) = \lambda_I(t_k \mid N_{1:k}, \theta_I)\lambda_E(t_k \mid N_{1:k}^{1:C}, \theta_E)\lambda_X(t_k \mid \mathbf{x}_{k+\tau}, \theta_X) \quad (6)$$

where $\theta = \{\theta_X, \theta_E, \theta_I\}$, $\lambda_I(t_k \mid N_{1:k}, \theta_I)$ is the component of the intensity function conditioned on the spiking history $N_{1:k}$ of the neuron whose intensity is being modeled, $\lambda_E(t_k \mid N_{1:K}^{1:C}, \theta_E)$ is the component related to the ensemble history contribution, and $\lambda_X(t_k \mid \mathbf{x}_{k+\tau}, \theta_X)$ is the component related to an extrinsic covariate $\mathbf{x}_{k+\tau}$, where τ is an integer time shift. Note that the term H_k, used in the previous section, is now replaced by more specific information according to the model.

We consider the following specific models for each of these 3 covariate types. We begin with a model incorporating the spiking history component.

The spiking history component is modeled as

$$\lambda_I(t_k \mid N_{1:k}, \theta_I) = \exp\left\{\gamma_0 + \sum_{n=1}^{Q} \gamma_n \Delta N_{k-n}\right\} \quad (7)$$

where Q is the order of the autoregressive process, γ_n represents the autoregressive coefficients, and γ_0 relates to a background level of activity. This model is henceforth referred to as the autoregressive spiking history model. We apply Akaike's standard information criterion (AIC, see *Eq. 16* below) to estimate the parameter Q. We expect this autoregressive spiking history model to capture mostly refractory effects, recovery periods, and oscillatory properties of the neuron.

The contributions from the ensemble are expressed in terms of a regression model of order R

$$\lambda_E(t_k \mid N_{1:k}^{1:C}, \theta_E) = \exp\left\{\beta_0 + \sum_{c} \sum_{r=1}^{R} \beta_r^c \Delta N_{k-r}^c\right\} \quad (8)$$

where the first summation is over the ensemble of cells with the exception of the cell whose conditional intensity function is being modeled. Thus the above model contains $R \times (C - 1)$ parameters plus one additional parameter for the background level. Note that the coefficients in the ensemble model capture spike effects at 1-ms time resolution and in this way they may reflect lagged-synchrony between spikes of the modeled cell and other cells in the ensemble. Alternatively, to investigate ensemble effects at lower time precision, we consider the ensemble rates model

$$\lambda_E(t_k \mid N_{1:k}^{1:C}, \theta_E) = \exp\left\{\beta_0 + \sum_{c} \sum_{r=1}^{R} \beta_r^c (N_{k-(r-1)W}^c - N_{k-rW}^c)\right\} \quad (9)$$

where the term $N_{k-(r-1)W}^c - N_{k-rW}^c$ is the spike count in a time window of length W covering the time interval $(t_{k-rW}, t_{k-(r-1)W}]$. The coefficients in this model may reflect spike covariances on slow time scales.

In our application to MI data, the extrinsic covariate $\mathbf{x}_{k+\tau}$ will specify the hand velocity. To model this component we employ a variation of the Moran and Schwartz (1999) model, henceforth referred to as the velocity model

$$\lambda_X(t_k \mid \mathbf{x}_{k+\tau}, \theta_X) = \exp\{\alpha_0 + |V_{k+\tau}|\,[\alpha_1 \cos(\phi_{k+\tau}) + \alpha_2 \sin(\phi_{k+\tau})]\} \quad (10)$$

where $|V_{k+\tau}|$ and $\phi_{k+\tau}$ are, respectively, the magnitude and angle of the 2-D hand velocity vector in polar coordinates at time $t_{k+\tau}$. In this model $\mathbf{x}_{k+\tau} = [|V_{k+\tau}|, \phi_{k+\tau}]'$. For illustration purposes, we have considered only a single, fixed-time shift τ in the above model. Based on previous results (Paninski et al. 2004) we set $\tau = 150$ ms. A much more generic model form including linear or nonlinear functions of covariates at many different time lags could be easily formulated.

The most complex conditional intensity function models we investigate are the autoregressive spiking history plus velocity and ensemble activity, and the autoregressive spiking history plus velocity and ensemble rates models. For the former, the full conditional intensity function model is given by

$$\lambda(t_k \mid N^{1:C}_{1:k}, \mathbf{x}_{k+\tau}, \theta) = \exp\{\mu + \sum_{n=1}^{Q} \gamma_n \Delta N_{k-n} + \sum_{c} \sum_{r=1}^{R} \beta^c_r \Delta N^c_{k-r} + |V_{k+\tau}|\,[\alpha_1 \cos(\phi_{k+\tau}) + \alpha_2 \sin(\phi_{k+\tau})]\} \quad (11)$$

where μ relates to the background activity.

It should be noticed that although these models are in the "generalized linear" model class, the relation between the conditional intensity function and spiking history, ensemble, and extrinsic covariates can be highly nonlinear. These models are linear only after the transformation of the natural parameter (here the conditional intensity function) by the log link function and only with respect to the model parameters being estimated. As seen in *Eq. 4*, general functions (e.g., quadratic, cubic, etc.) of the actual measured covariates can be used.

NON-GLM CONDITIONAL INTENSITY FUNCTION MODEL. To illustrate the generality of the proposed point process framework, we construct and analyze a non-GLM conditional intensity function model that also incorporates effects of spiking history, neural ensemble, and extrinsic covariates. Additionally, this example demonstrates a procedure for obtaining a conditional intensity function by first modeling the interspike interval (ISI) conditional probability density function. The conditional intensity is obtained from the ISI probability density model using the relation (Brown et al. 2003)

$$\lambda(t_k \mid \theta, H_k) = \frac{p(t_e \mid \theta, H_k)}{1 - \int_{u_{N_{k-1}}}^{t_k} p[t \mid \theta, H(t)]\mathrm{d}t} \quad (12)$$

where $t_e = t_k - u_{N_{k-1}}$ is the elapsed time since the most recent spike of the modeled cell before time t_k and $p(t_e \mid \theta, H_k)$ is the ISI probability density, specified here by the inhomogeneous inverse Gaussian (Barbieri et al. 2001). This probability density is given in *Eq. A11* in the APPENDIX. This density is specified by a time-varying scaling parameter $s(t_k \mid \cdot)$ that, in our application to MI spiking data, captures the velocity and ensemble rates covariate effects

$$s(t_k \mid \mathbf{x}_{t+\tau}, N^{1:C}_{1:k}, \theta_X, \theta_E) = \exp\left\{\mu + \sum_{c} \sum_{r=1}^{R} \beta^c_r (N^c_{k-(r-1)W} - N^c_{k-rW}) + |V_{k+\tau}|\,[\alpha_1 \cos(\phi_{k+\tau}) + \alpha_2 \sin(\phi_{k+\tau})]\right\} \quad (13)$$

and a location parameter ψ. The set of parameters defining the inhomogeneous inverse Gaussian density and therefore the conditional intensity function is denoted $\theta = \{\theta_X, \theta_E, \psi\}$. This model (*Eqs. 12*, *13*, and *A11*) is henceforth referred to as the inhomogeneous inverse Gaussian plus velocity and ensemble rates model. The history dependence in this model extends back to the time of the previous, most recent spike.

Maximum likelihood parameter estimation

Maximum likelihood parameter estimates for the models in the point process–GLM framework were efficiently computed using the iterative reweighted least squares (IRLS) algorithm. This method is the standard choice for the maximum likelihood estimation of GLMs because of its computational simplicity, efficiency, and robustness. IRLS applies the Newton–Raphson method to the reweighted least squares problem (McCullagh and Nelder 1989). Given the conditional intensity model in *Eq. 4*, the log-likelihood function is strictly concave. Therefore, if the maximum log-likelihood exists, it is unique (Santner and Duffy 1989). Confidence intervals and *p*-values were obtained following standard computations based on the observed Fisher information matrix (Pawitan 2001). Statistically nonsignificant parameters (e.g. $P \geq 0.001$) were set to zero for all of the models. In the non-GLM case, the inhomogeneous inverse Gaussian model was fit by direct maximization of the likelihood function using a quasi-Newton method (IMSL, C function *min_uncon_multivar*, from Visual Numerics, 2001). For the data sets used here, the most intensive computations involved operations on large matrices of size about $10^6 \times 200$. Algorithms were coded in C and run on dual-processor 3.9-GHz IBM machines, 2 GB of RAM memory. Standard GLM estimation using IRLS is also available in several statistical packages (S-Plus, SPSS, and Matlab Statistics toolbox).

Goodness-of-fit, point process residual analyses and model comparison

KOLMOGOROV–SMIRNOV (K-S) TEST ON TIME RESCALED ISIS. Before making an inference from a statistical model, it is crucial to assess the extent to which the model describes the data. Measuring quantitatively the agreement between a proposed model and a spike train data series is a more challenging problem than for models of continuous-valued processes. Standard distance measures applied in continuous data analyses, such as average sum of squared errors, are not designed for point process data. One alternative solution to this problem is to apply the time-rescaling theorem (Brown et al. 2002; Ogata 1988; Papangelou 1972) to transform point processes like spike trains into continuous measures appropriate for goodness-of-fit assessment. Once a conditional intensity function model has been fit to a spike train data series, we can compute rescaled times z_j from the estimated conditional intensity and from the spike times as

$$z_j = 1 - \exp\left\{-\int_{u_j}^{u_{j+1}} \lambda(t \mid H(t), \hat{\theta})\mathrm{d}t\right\} \quad (14)$$

for $j = 1, \ldots, J-1$, where $\hat{\theta}$ is the maximum likelihood estimate of the parameters. The z_j values will be independent uniformly distributed random variables on the interval [0, 1) if and only if the conditional intensity function model corresponds to the true conditional intensity of the process. Because the transformation in *Eq. 14* is one to one, any statistical assessment that measures the agreement between the z_j values and a uniform distribution directly evaluates how well the original model agrees with the spike train data. To construct the K-S test, we order the z_j values from smallest to largest, denoting the ordered values as $z_{(j)}$, and then plot the values of the cumulative distribution function of the uniform density function defined as $b_j = (j - 1/2)/J$ for $j = 1, \ldots, J$ against the $z_{(j)}$. We term these plots *K-S plots*. If the model is correct, then the points should lie on a 45° line. Confidence bounds for the degree of agreement between a model and the data may be constructed using the distribution of the Kolmogorov–Smirnov statistic (Johnson and Kotz 1970). For moderate to large sample sizes the 95% confidence bounds are well approximated by $b_j \pm 1.36 \cdot J^{-1/2}$ (Johnson and Kotz 1970).

To assess how well a model performs in terms of the original ISIs ($ISI_j = u_j - u_{j-1}$), we relate the ISIs to the computed z_j values in the following manner. First, the empirical probability density of the z_j values is computed, and the ratio of the empirical to the expected (uniform) density is calculated for each bin in the density. Second, the ISI values in the data are rounded to integer milliseconds and collected into bins. For these ISIs, all the corresponding z_j values as well as the ratios of empirical to expected densities in the related bins are obtained. This correspondence between ISIs and z_j values is easily available from *Eq. 14*. Third, we compute the mean ratio R (i.e., the mean of all the ratios for this particular ISI value). A mean ratio $R > 1$ ($R < 1$) implies that there are more (less) rescaled ISIs of length z_j than expected and that the intensity is being underestimated (overestimated), on average, at this particular ISI value.

If the model is correct, the z_j values should be not only uniformly distributed, but also independent. Thus, even when the K-S statistic is small, we still need to show that the rescaled times are independent. Here we assess independence up to 2nd-order temporal correlations by computing the autocorrelation function of the transformed rescaled times. As a visualization aid, we plot z_{j+1} against z_j.

POINT PROCESS RESIDUAL ANALYSIS. A standard approach in goodness-of-fit analysis is to examine structure in the data that is not described by the model. For continuous valued data, this is done by analyzing the residual error (i.e., the difference between the true and predicted values). For point process data, a different definition of residuals is needed to relate the conditional intensity function to the observed spike train data. The point process residual (Andersen et al. 1992) over nonoverlapping moving time windows is defined as

$$M(t_k) = \sum_{i=k-B}^{k} \Delta N_i - \int_{t_{k-B}}^{t_k} \lambda(t \mid H(t), \hat{\theta})\mathrm{d}t \tag{15}$$

for $k - B \geq 1$. In the application to MI data, we will look for relations between the point process residual and motor covariates (e.g., speed or direction) by computing their cross-correlation functions. Existence of correlations would imply that there is some structure left in the residuals that is not captured by the conditional intensity function model.

MODEL SELECTION. An additional tool for comparing models comes from the statistical theory of model selection (Burnham and Anderson 2002). The idea consists of choosing the best models to approximate an underlying process generating the observed data, a process whose complexity can be potentially infinite dimensional. To achieve this goal, we adopt Akaike's standard information criterion (AIC) (Akaike 1973). This criterion also provides a way to rank different candidate models. The AIC was originally derived as an estimate of the expected relative Kullback–Leibler distance (Cover and Thomas 1991) between a distribution given by an approximating model and the distribution of the true underlying process generating the data. This criterion is formulated as

$$AIC(q) = -2 \log L(\hat{\theta} \mid H_k) + 2q \tag{16}$$

where $L(\hat{\theta} \mid H_K)$ is the likelihood function, $L(\hat{\theta} \mid H_K) = P(N_{1:K} \mid \hat{\theta}, H_K)$; $\hat{\theta}$ is the maximum likelihood estimate of the model parameters $\hat{\theta}$; and q is the total number of parameters in the model. By this criterion, the best model is the one with the smallest AIC, implying that the approximate distance between this model and the "true process" generating the data is smallest. The AIC is frequently interpreted as a measure of the trade-off between how well the model fits the data and the number of parameters required to achieve that fit, or of the desired trade-off between bias and variance (Burnham and Anderson 2002). An equivalence between AIC and cross-validation for the purpose of model selection has been established (Stone 1977). AIC can be applied to both nested and nonnested models, and to models with different distributions in their stochastic component. We compute AIC values for different models to guide our model comparison. Specifically, we provide the difference between the AIC of all of the models with respect to the AIC of the best model. We also use the AIC to estimate the order of the autoregressive spiking history component in *Eq. 7*.

Neural decoding analysis by state estimation with point process observations

Beyond assessing the goodness-of-fit of a single cell model with respect to its individual spike train data, we also analyze the ability of the model, over the entire cell population, to decode an m-dimensional extrinsic covariate $\mathbf{x}_{k+\tau}$. Such decoding will use the spike times of the entire ensemble of cells and the corresponding conditional intensity function for each of these cells. We thus perform a state estimation of $\mathbf{x}_k$ based on point process observations and thereby assess the ensemble coding properties of the cell population. The estimated extrinsic covariate will be given by the posterior mode after a Gaussian approximation to the Bayes–Chapman–Kolmogorov system (Eden et al. 2004).

For the particular type of hand kinematics data described above, we model $\mathbf{x}_k$ as a Gaussian autoregressive process of order 1, henceforth $AR(1)$, given by

$$\mathbf{x}_{k+\tau} = \mu_{\mathbf{x}} + \mathbf{F}\mathbf{x}_{k+\tau-1} + \varepsilon_{k+\tau} \tag{17}$$

where $\mu_{\mathbf{x}}$ is an m-dimensional vector of mean parameters, $\mathbf{F}$ is an $m \times m$ state matrix, and ε_k is the noise term given by a zero mean m-dimensional white Gaussian vector with $m \times m$ covariance matrix $\mathbf{W}_{\varepsilon}$. The matrices $\mathbf{F}$ and $\mathbf{W}_{\varepsilon}$ are fitted by maximum likelihood.

The point process observation equation is expressed in terms of the modeled conditional intensity functions $\lambda^c(t_k \mid \cdot)$ for each of the C cells entering the decoding. As an example, for intensity functions conditioned on a motor covariate $\mathbf{x}_{k+\tau}$ and intrinsic spiking history $N^C_{1:k}$, we have the following recursive point process filter.

One step prediction

$$\mathbf{x}_{k+\tau \mid k+\tau-1} = \mu_{\mathbf{x}} + \mathbf{F}\mathbf{x}_{k+\tau-1 \mid k+\tau-1} \tag{18}$$

One-step prediction covariance

$$\mathbf{W}_{k+\tau \mid k+\tau-1} = \mathbf{F}\mathbf{W}_{k+\tau-1 \mid k+\tau-1}\mathbf{F} + \mathbf{W}_{\varepsilon} \tag{19}$$

Posterior covariance

$$\mathbf{W}_{k+\tau \mid k+\tau} = \Bigg[\mathbf{W}^{-1}_{k+\tau \mid k+\tau-1} + \sum_{c=1}^{C} [\nabla \log \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c)] \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c)\Delta [\nabla \log \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c)]' - \sum_{c=1}^{C} \nabla^2 \log \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c) [\Delta N^c_{1:k} - \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c)\Delta] \Bigg]^{-1} \tag{20}$$

Posterior mode

$$\mathbf{x}_{k+\tau \mid k+\tau} = \mathbf{x}_{k+\tau \mid k+\tau-1} + \mathbf{W}_{k+\tau \mid k+\tau} \times \sum_{c=1}^{C} \nabla \log \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c) [\Delta N^c_{1:k} - \lambda^c(t_k \mid N^c_{1:k}, \mathbf{x}_{k+\tau \mid k+\tau-1}, \hat{\theta}^c)\Delta] \tag{21}$$

The term $\nabla(\nabla^2)$ denotes the m-dimensional vector ($m \times m$ matrix) of first (second) partial derivatives with respect to $\mathbf{x}_{k+\tau}$, and $\mathbf{W}_{k+\tau \mid k+\tau}$ is the posterior covariance matrix of $\mathbf{x}_{k+\tau}$. Similarly, decoding equations based on other models of the conditional intensity function can be obtained. The derivation of the recursive point process filter is based on the well-established (Mendel 1995; Kitagawa and Gersh 1996) relation between the posterior probability density and the Chapman–Kolmogorov (one-step prediction) probability density, and on a Gaussian approximation of the posterior density (for details see Eden et al. 2004). The Gaussian approximation results from a 2nd-order Taylor expansion of the density and it is a standard first approach for approximating probability densities (Tanner 1996; Pawitan 2001). Nonetheless, the spiking activity enters into the computations in a very non-Gaussian way through the point process model instantiated by the conditional intensity function.

The amount of uncertainty in the algorithm about the true state of the decoded parameter is related to the matrix $\mathbf{W}_{k+\tau \mid k+\tau}$. Confidence

regions and coverage probability for the decoding can thus be obtained as follows. At time $k\Delta t$ an approximate 0.95 confidence region for the true covariate $\mathbf{x}_{k+\tau}$ may be constructed as

$$(\mathbf{x}_{k+\tau} - \mathbf{x}_{k+\tau|k+\tau})' \mathbf{W}^{-1}_{k+\tau|k+\tau} (\mathbf{x}_{k+\tau} - \mathbf{x}_{k+\tau|k+\tau}) \leq \chi^2_{0.95}(m) \quad (22)$$

for $k = 1, 2, \ldots, K$, where $\chi^2_{0.95}(m)$ gives the 0.95 quantile of the χ^2 distribution with degrees of freedom equal to the dimension m of the covariate. The coverage probability up to time t_k is given by s_k/k where s_k is the number of times the true covariate is within the confidence regions during the time interval $(0, k\Delta]$. In the decoding analysis we compute the mean of the coverage probability over the entire decoding period. A Monte Carlo simulation is employed to obtain the confidence intervals and coverage probability for the covariate in polar coordinates. We first use the estimated posterior covariance matrix to generate 10^4 Gaussian-distributed samples centered at the current covariate estimates in Cartesian coordinates. Second, these random samples are converted to polar coordinates. Finally, the confidence intervals are then computed from the distribution of the random samples in polar coordinates.

RESULTS

The proposed point process framework is illustrated with 2 examples. The first one is applied to simulated neural spike data and the second to multiple single units simultaneously recorded from monkey primary motor cortex. For the discrete time representation of the neural point process we set $\Delta = 1$ ms.

Simulation study

The goal of the simulation study is 2-fold. First, we illustrate the main properties of the model in *Eq. 11* containing the autoregressive history, neural ensemble history, and motor covariate effects. Second, we demonstrate that the parameters of the simulated model are accurately recovered from relatively small spike data sets by maximum likelihood estimation implemented with the IRLS algorithm.

The conditional intensity functions of 6 neurons (A, B, C, D, E, F) were simulated using methods as described in Ogata (1981). The intensity of 5 of them (B–F) was given by the velocity model (*Eq. 10*); that is, the neurons were modeled as inhomogeneous Poisson processes with mean background spiking rates of 17, 16, 9, 8, and 7 Hz, respectively, and inhomogeneity introduced by the modulating hand velocity signal. Different velocity tuning functions were used for the set of cells. The hand velocity signal was sampled from actual hand trajectories performed by a monkey (see *Application to MI spiking data*, below). The conditional intensity function for neuron A was given by the autoregressive spiking history plus ensemble and velocity model (*Eq. 11*). The background mean firing rate of this neuron was set to 10 Hz. The autoregressive spiking history component contained 120 coefficients covering 120 ms of spiking history (see Fig. 2*B*). The autoregressive coefficients mimicked the effects of refractory–recovery periods and rebound excitation. From the ensemble of 5 neurons, only 2 contributed excitatory (neuron B) and inhibitory (neuron C) effects at 3 time lags (-1, -2, and -3 ms).

The simulation scheme worked as follows. Starting with the initial simulation time step, first the conditional intensity functions were updated and then, at the same time step, the spiking activities for all of the cells were simulated. The simulation then moved to the next time step. The conditional intensity functions were updated based on the past intrinsic and ensemble spiking history (neuron A only) and on the current hand velocity state (all neurons).

The main features of the conditional intensity function model in *Eq. 11* can be observed in Fig. 1, where the simulated conditional intensity function of neuron A and its own spiking activity are plotted together with the activity of the other 5 neurons and the contribution of velocity signal. The simulated conditional intensity function clearly shows the dependence on spike history: after a spike, the intensity drops to almost zero and slowly recovers, reaching a period of higher than background spiking probability at about 20 ms after the spike. Fast excitatory and inhibitory effects follow the spikes of neurons B and C. Spiking history, neural ensemble, and velocity modulate each other's contributions in a multiplicative fashion.

From the simulated ensemble spike trains and from the velocity signal, we then estimated the conditional intensity function generating the spiking activity of neuron A. The data set entering the estimation algorithm thus consisted of 6 simulated spike trains, each 200 s long, and of the hand velocity time series in polar coordinates. The spike train for the modeled neuron A contained 2,867 spikes. The parametric model for the estimated conditional intensity function consisted of a background mean, 120 autoregressive coefficients, and 5 regressive coefficients for each of the other 5 neurons (i.e., 25 coefficients in total). Each set of 5 coefficients related to spiking activity at lags $-1, -2, \ldots, -5$ ms. The IRLS

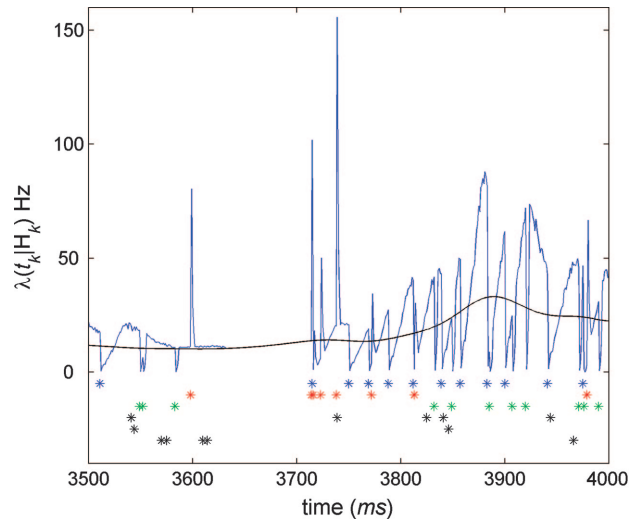

FIG. 1. Simulated conditional intensity function model. Conditional intensity function, modeled as in *Eq. 1*, was simulated and used to generate a spike train (neuron A, blue asterisks mark the times of the spike events). In this model, the intensity (blue curve) was conditioned on the past spiking history, the spikes of 2 other neurons (neuron B, excitatory, red asterisks; neuron C, inhibitory, green asterisks), and on hand velocity. Past spiking history effect was modeled by a 120-order autoregressive process carrying a refractory period, recovery, and rebound excitation. Coefficient values were based on parameters estimated from a primary motor cortex (MI) cell (see Fig. 5*B*). The conditional intensity function resulting from the contribution of only hand velocity is shown by the black line. Three other cells were also simulated (neurons D, E, and F; black asterisks). Neurons B–F were modeled as inhomogeneous Poisson processes modulated according to the velocity model (*Eq. 10*). All cells had different preferred movement directions. Spiking history, ensemble, and velocity modulated each other in a multiplicative fashion. Simulated ensemble spike trains together with hand velocity were used to estimate the parameters for the conditional intensity function model of neuron A (see Fig. 2).

algorithm converged in 12 iterations (tolerance was set to 10^{-6}). Statistically nonsignificant parameters ($P > 0.001$) were set to zero (see METHODS section). The true model parameters used in the simulation of neuron A were accurately recovered (Fig. 2, *B* and *C*), with the estimated model passing the K-S goodness-of-fit test (Fig. 2*D*). Parameter estimation on smaller data sets (about 50 s of data) led to similar successful fittings.

Application to MI spiking data

Experimental data were obtained from the MI area of a behaving monkey. Details of the basic recording hardware and protocols are available elsewhere (Donoghue et al. 1998; Maynard et al. 1999). After task training, a Bionic Technologies LLC (BTL, Salt Lake City, UT) 100-electrode silicon array was implanted in the area of MI corresponding to the arm representation. One monkey (*M. mulatta*) was operantly conditioned to track a smoothly and randomly moving visual target. The monkey viewed a computer monitor and gripped a 2-link, low-friction manipulandum that constrained hand movement to a horizontal plane. The hand (*x, y*) position signal was digitized and resampled to 1 kHz. Low-pass–filtered finite differences of position data were used to obtain hand velocities. Some 130 trials (8–9 s each) were recorded. More details about the statistical properties of the distributions for hand position and velocity, spiking sorting methods, and other task details can be found in Paninski et al. (2004).

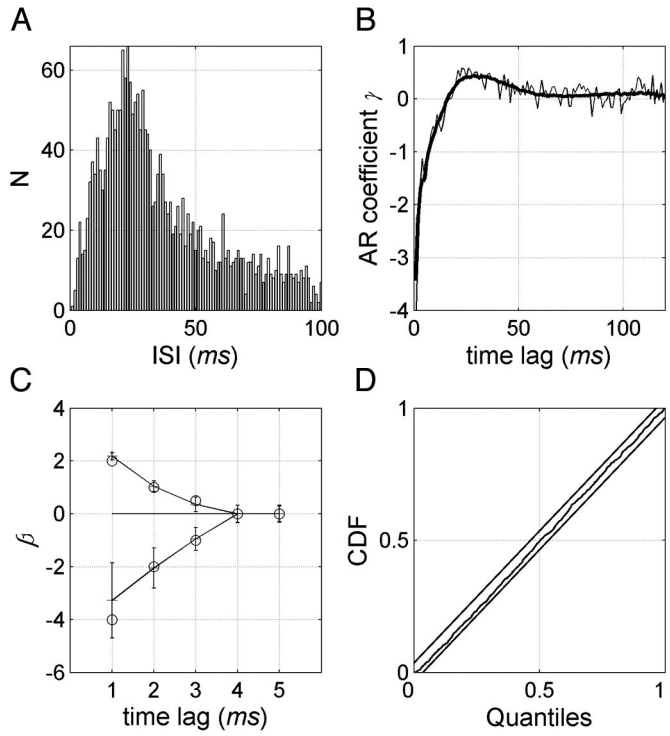

FIG. 2. Model parameter estimation by iteratively reweighted least squares (IRLS) in the generalized linear model (GLM) framework. Spike trains of the 6 simulated cells (each lasting 200 s; see Fig. 1) together with hand velocity produced the data set for the estimation of the parameters of the conditional intensity function model for neuron A. Spike train of neuron A consisted of 2,867 spikes. *A*: ISI distribution for neuron A. *B*: true autoregressive coefficients (thick curve) and the estimated ones. *C*: ensemble covariate in the estimated model contained 5 coefficients per cell (covering 5 ms of the past). Only the 3 coefficients for the excitatory and inhibitory cells were significantly different from zero ($P < 0.001$). Bars indicate the 95% confidence intervals. Small circles show the location of the true coefficients. *D*: Kolmogorov–Smirnov (K-S) goodness-of-fit test shows that the estimated conditional intensity model passed the test (the 95% confidence region is given by the parallel lines). Coefficients for the velocity covariate $\alpha_1 = 0.1$ and $\alpha_2 = -0.05$ were estimated as $\hat{\alpha}_1 = 0.105$ and $\hat{\alpha}_2 = -0.045$.

Models including 1, 2, or all of the 3 types of covariates were analyzed. To start, we focus on the analysis of the velocity and the autoregressive spiking history plus velocity models. Later, we also compare these 2 models using neural decoding based on the observation of the entire ensemble of cells. For this reason, we analyzed these 2 models for each of the 20 cells in the ensemble. More detailed analysis involving K-S plots, point process residuals, and AIC model comparison will be illustrated for one typical cell.

K-S GOODNESS-OF-FIT ANALYSIS FOR THE VELOCITY AND THE AUTOREGRESSIVE SPIKING HISTORY PLUS VELOCITY MODELS. The tuning functions obtained from the velocity model (*Eq. 10*) are shown in Fig. 3. This model was statistically significant for all of the cells. Preferred direction was diverse across cells, covering the range of possible directions. The corresponding K-S plots are shown in Fig. 4. The quantiles refer to the $z_{(j)}$s (*Eq. 14*) and the cumulative distribution function (CDF) refers to the expected uniform distribution for the case when the estimated conditional intensity model was equivalent to the true one. The velocity model tends to overestimate (underestimate) the conditional intensity at lower (middle) quantiles. Introduction of the autoregressive spiking history component (*Eq. 7*) in the velocity model greatly improved the explanation of the spiking process, almost completely eliminating both the over- and underestimation of the intensity. The maximum order of the autoregressive component was about 120 (i.e., the component incorporated history effects spanning over 120 ms in the past). The most significant history effects extended to 60 ms in the past. For the majority of the cells, this component seemed to capture mostly 3 main history effects: refractory and recovery periods followed by an increase in the firing probability around 20 ms after a spike (see Fig. 5*B*). It should be noticed that the autoregressive coefficients could have also reflected dynamical network properties of nonmeasured neural ensembles such as networks of excitatory and inhibitory neurons where the modeled cell is embedded, or nonmeasured fast extrinsic covariates. No significant differences in the K-S plots were observed between a pure autoregressive history model and the autoregressive history plus velocity models (not shown).

Figure 5, *C* and *D* summarize the above observations for a typical cell (cell 75a, 29,971 spikes over 130 trials, i.e., ~1,040 s) in this example set and relate the fitting problems of the velocity model to the original nontime rescaled ISIs. In the velocity model, the intensity is overestimated (mean ratio $R < 1$) for ISIs below 10 ms and underestimated (mean ratio $R > 1$) for ISIs in the interval 10 to about 40 ms (Fig. 5*D*). The overestimation is likely a reflection of a refractory–recovery period (up to ~10 ms) after the cell has spiked, which is not captured by the velocity model. The underestimation reflects a period of increased firing probability that follows the recovery period. These 2 different regimes are reasonably well captured by the coefficients of the autoregressive component (see Fig. 5*B*), thus resulting in the improved fit observed for the autoregressive spiking history plus velocity model. Introduction of

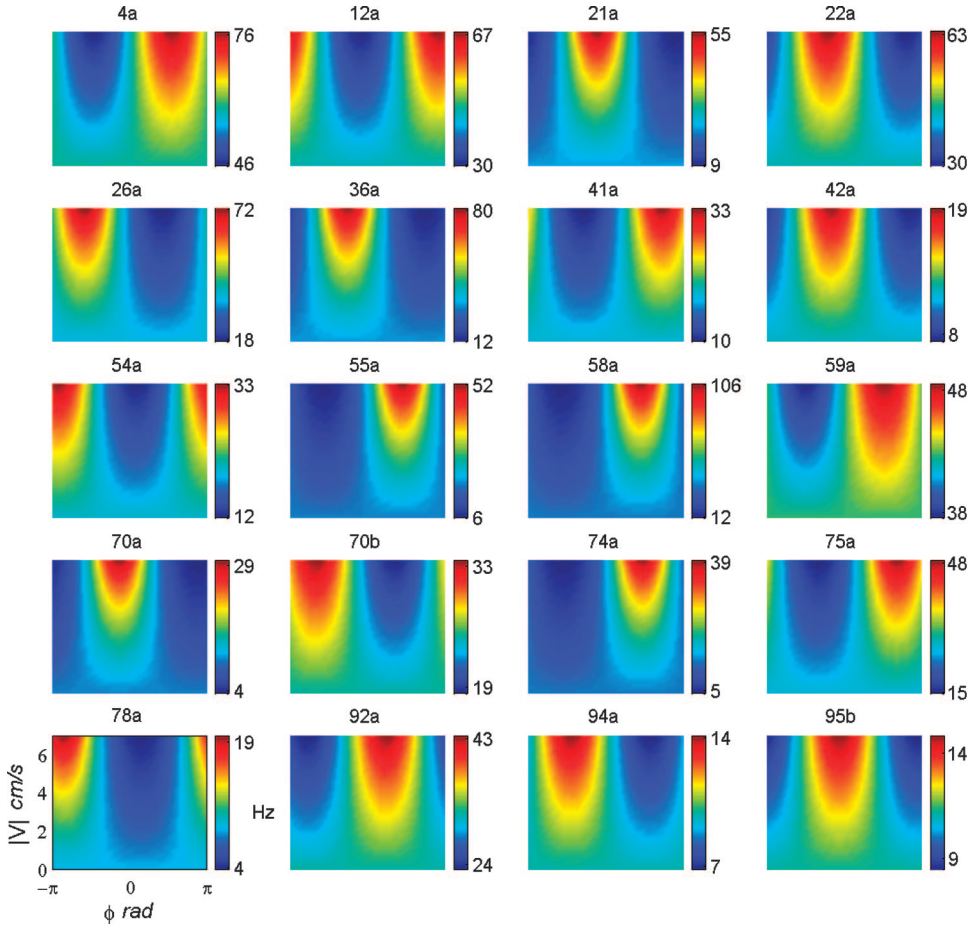

FIG. 3. Velocity tuning functions. Conditional intensity function values, based on the velocity model, are expressed by pseudocolor maps. Velocity is given in polar coordinates, with ϕ representing the movement direction. Each subplot relates to a particular cell, with cells' labels given at the *top*.

the autoregressive component makes the observed density for the z_j values much closer to the expected uniform density (Fig. 5*C*).

The K-S statistic measures how close rescaled times are to being uniformly distributed on [0, 1). In addition, a good model should also generate independent and identically distributed rescaled times. To illustrate this point, we checked for temporal correlations at lag 1 in the time-rescaled ISIs (Fig. 6). As expected, some temporal structure remains in the case of the velocity model ($r^2 = 0.25$, $P = 10^{-6}$), whereas this structure is effectively insignificant for the velocity plus autoregressive spiking history ($r^2 = 0.002$, $P = 10^{-6}$). The cross-correlation function computed over a broad range of lags was consistent with this result.

POINT PROCESS RESIDUAL ANALYSIS. Even though the parameters for the velocity model were statistically significant, the K-S plot analysis showed that the velocity model fell short of explaining the entire statistical structure in the observed single-cell spiking activity. It thus remains to be seen how well this model captured the relationship between hand velocity and spiking activity. Besides neural decoding, another approach to address this problem is to measure the correlations among the point process residuals as defined in *Eq. 15* and the movement velocity. Existence of correlations would imply that there is some structure left in the residuals that is not captured by the velocity model. On the other hand, a decrease in the correlation level with respect to some other model would imply that the velocity model does capture some of the structure in the spiking activity related to hand velocity.

We computed the correlations for the residuals from the autoregressive spiking history model (*Eq. 7*) and compared them to the correlations for the residuals from the velocity and from the autoregressive spiking history plus velocity model. Residuals were computed for nonoverlapping 200-ms moving windows (Fig. 7). Cross-correlation functions were computed between the residuals and the mean of the kinematic variables. Mean (x, y) velocities were computed for each time window and were used to obtain, in polar coordinates, the respective mean movement speed and direction. In the autoregressive model case, peak cross-correlation values between the residuals and direction, speed, and velocities in x and y coordinates were 0.29, 0.10, −0.17, and 0.50, respectively. For the autoregressive spiking history and velocity model, the peak cross-correlation values for the same variables were 0.08, 0.06, −0.12, and 0.28. This suggests that, for this particular neuron, the velocity model captures a significant amount of information about hand velocity available in the spiking activity. Nonetheless, it is also clear that there is a residual structure in the spiking activity that is statistically related to the hand velocity in Cartesian coordinates and that is not captured by the autoregressive spiking history plus velocity model. Furthermore, the cross-correlation functions for both the velocity and the autoregressive spiking history plus velocity model show no significant differences, which suggests that the autoregressive

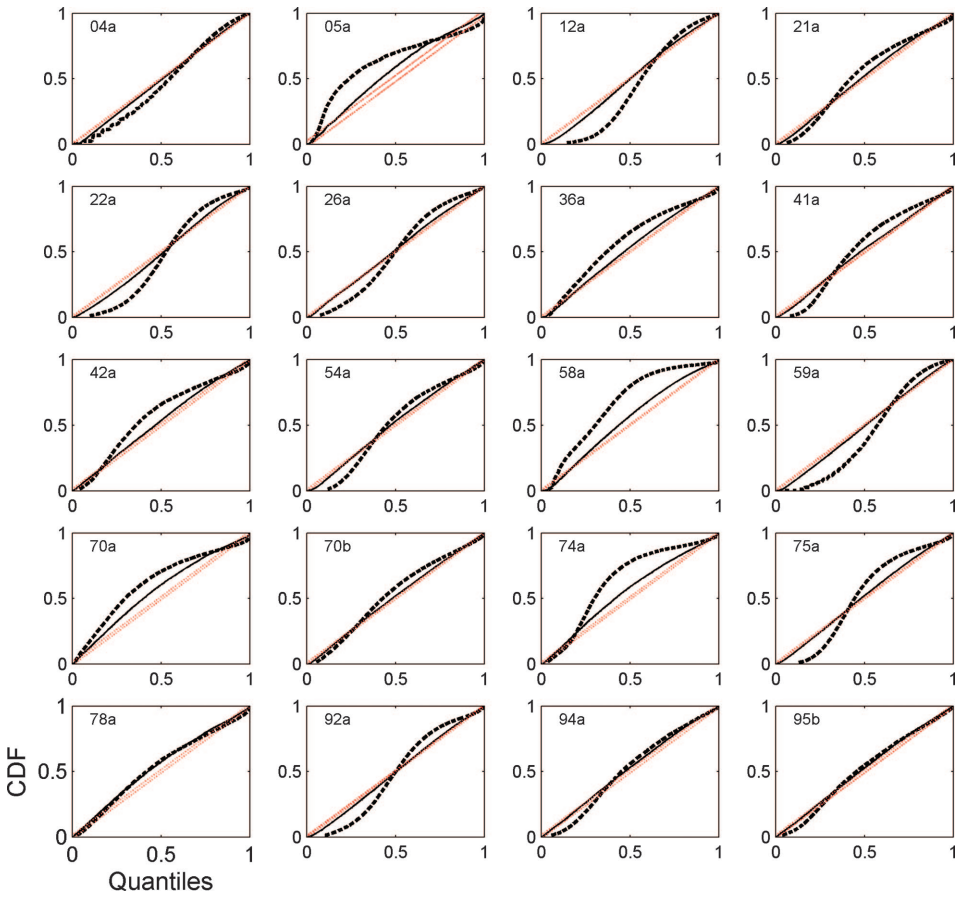

FIG. 4. K-S plots for the velocity model and the autoregressive spiking history plus velocity model. Because the K-S plots are constructed from a uniform distribution on the interval [0, 1), the 50th and 100th percentiles correspond, respectively, to quantiles 0.5 and 1.0 on both the horizontal and vertical axes. Two-sided 95% confidence error bounds of the K-S statistics are displayed for each cell (45° red lines). Visual inspection alone already reveals that, for most of the cells, the autoregressive spiking history plus velocity model (solid curve) improves the fit considerably.

component does not carry additional statistical information about hand velocity.

MODEL COMPARISON. We compared partial and complete (i.e., 1, 2, or 3 covariate types) conditional intensity function models for cell 75a. The ensemble model included spiking activity of each cell at 4 time lags (−1, −2, −3, and −4 ms). The ensemble rates included the spike counts of each cell at 3 nonoverlapping and lagged time windows. The length of each of the time windows, specified by the parameter W in *Eq. 9*, was 50 ms. The K-S plots in Fig. 8 reveal that the autoregressive spiking history plus velocity and the autoregressive spiking history plus velocity and ensemble rates models provided the best fits among all of the models for this specific cell. The inhomogeneous inverse Gaussian plus velocity and ensemble rates model performed better than the velocity, ensemble, and ensemble rates models. Inspection of the coefficients for the ensemble and ensemble rates models showed that the dependencies were statistically significant for many of the cells in the ensemble. Individual cells contributed either positive or negative effects to the conditional intensity function and the effective ensemble contribution to the modulation of the conditional intensity function could reach tens of hertz.

In the above K-S plot comparisons, some of the models had K-S statistics far from the 95% confidence intervals or nearly identical to those from other models, making a clear comparison difficult. The AIC analysis was then used to provide a more detailed comparison, as well as to take the complexity of the model (i.e., number of parameters) into consideration in the model comparison. Figure 9 shows the ranked models in terms of their difference with respect to the AIC of the best model. In this context, models with lower AIC difference values are considered better models.

Overall, this criterion provided a fine model ranking and suggested that models containing the autoregressive spiking history component performed better in each instance. Among the alternative models for spiking history, the autoregressive spiking history model performed better than the conditional intensity model based on the inhomogeneous inverse Gaussian ISI distribution model (*Eqs. 12*, *13*, and *A11*), both in the AIC and K-S goodness-of-fit analyses. Also, the ensemble rates model did better than models containing only the velocity covariate or the ensemble covariate at fine temporal precision.

VELOCITY AND MOVEMENT DIRECTION DECODING ANALYSIS. The velocity (*Eq. 10*) and the autoregressive spiking history plus velocity models were used in the neural decoding of hand velocity. Models were fit to a training data set (120 trials, about 8–9 s each) and applied to decoding on a different test data set (10 trials, again about 8–9 s each). The state matrix $\mathbf{F}$ for the $AR(1)$ state process (*Eq. 17*) was estimated to be diagonal with nonzero terms approximately equal to 0.99, and the noise covariance matrix $\mathbf{W}_{\varepsilon}$ to be diagonal with nonzero entries equal to 0.01. Figure 10 shows the resulting decoding of movement direction and, in Cartesian coordinates, the estimated (x, y) velocities for a single test trial based on the velocity model. Overall, decoding of movement direction was remarkably good. Decoded (x, y) velocities captured mostly slower fluctuations. To compare the decoding performance of the 2 models, we computed the coverage probability and the

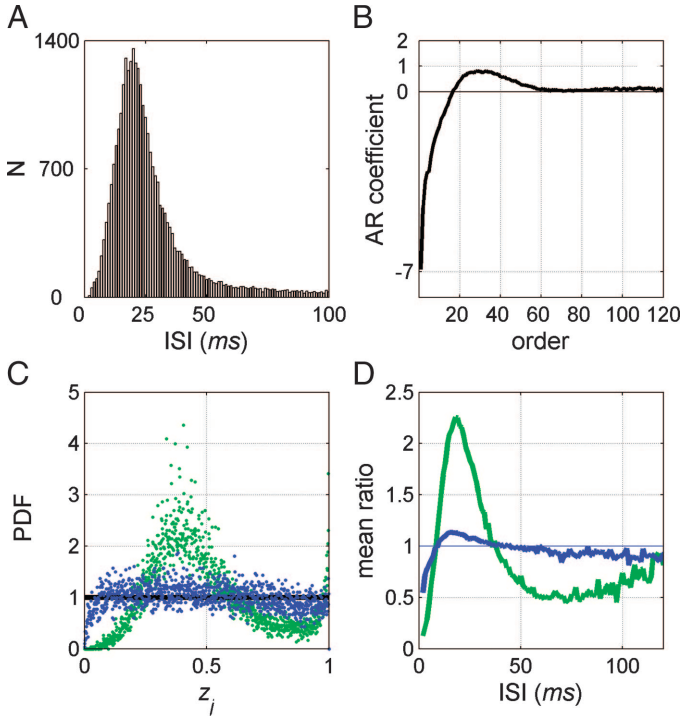

FIG. 5. Contribution of the autoregressive spiking history component. Cell 75a is chosen to illustrate how the addition of the autoregressive component improves the model's fit. *A*: ISI histogram. *B*: estimated coefficients of the autoregressive component. Autoregressive component incorporates a recovery period after the cell spikes, which lasts for about 18 ms (negative coefficients). Cell's firing probability then starts to increase, peaking at about 25 ms after a spike. *Order* refers to the order of the AR coefficient representing increasing times since the last spike. *C*: histogram for the transformed times z_j for both models (green: velocity model; blue: autoregressive spiking history plus velocity model). Black line shows the expected uniform distribution for the case where the estimated intensity function is close enough to the true intensity function underlying the neural point process. *D*: mean ratio of observed to expected z_j values indicates that the velocity model overestimates, on average, the intensity function for periods up to about 10 ms after a spike, while it tends to underestimate the intensity for periods between 10 and 40 ms. Introduction of the negative (positive) autoregressive coefficients almost completely eliminates the over (under) estimation of the conditional intensity function based on the velocity model alone.

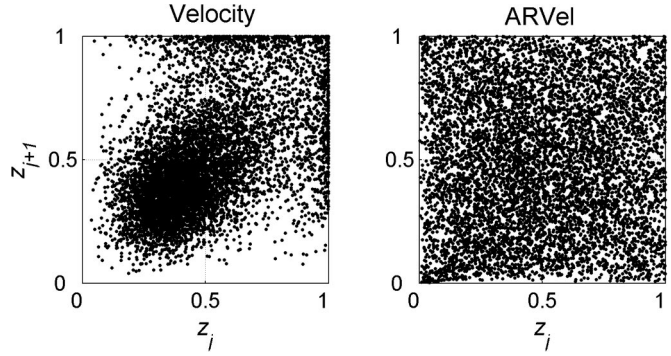

FIG. 6. Temporal correlations in the time-rescaled ISIs. Scatter plots are shown for consecutive z_j values from the velocity and autoregressive spiking history plus velocity (ARVel) models applied to cell 75a. Clearly, the autoregressive spiking history plus velocity model presents a more independent rescaled distribution. Corresponding correlation coefficients are 0.25 ($P = 10^{-6}$) for the velocity model and 0.002 ($P = 10^{-6}$) for the autoregressive spiking history plus velocity model. Cross-correlation functions computed over a broad range of lags led to similar results. Thus, in addition to improving the fit in the K-S plots, the introduction of the autoregressive component also eliminates temporal correlations among the rescaled times observed for the velocity model.

decoding error. Table 1 gives the mean values (across time and trials) for the coverage probabilities of the bivariate estimate (velocity magnitude and movement direction) and the coverage probabilities of the univariate estimate (velocity magnitude or movement direction). Mean coverage probability for the movement direction estimate was 0.94 for the velocity model. For the same model, coverage probabilities for the bivariate estimate and velocity magnitude were much smaller, consistent with the observation that the estimated velocities in Cartesian coordinates captured mostly slow fluctuations. Mean coverage probability, mean and median decoding errors, and confidence intervals for the decoding errors were not significantly different between the 2 models.

DISCUSSION

An important problem in neurophysiology is determining the factors that affect a neuron's spiking behavior. To address this question we have presented a point process statistical framework that allowed us to characterize simultaneously the effects of several covariates on the spiking activity of an individual neuron. The 3 types of covariates we considered were the neuron's spiking history, past neural ensemble activity, and extrinsic covariates such as stimuli or behavior. Because de-

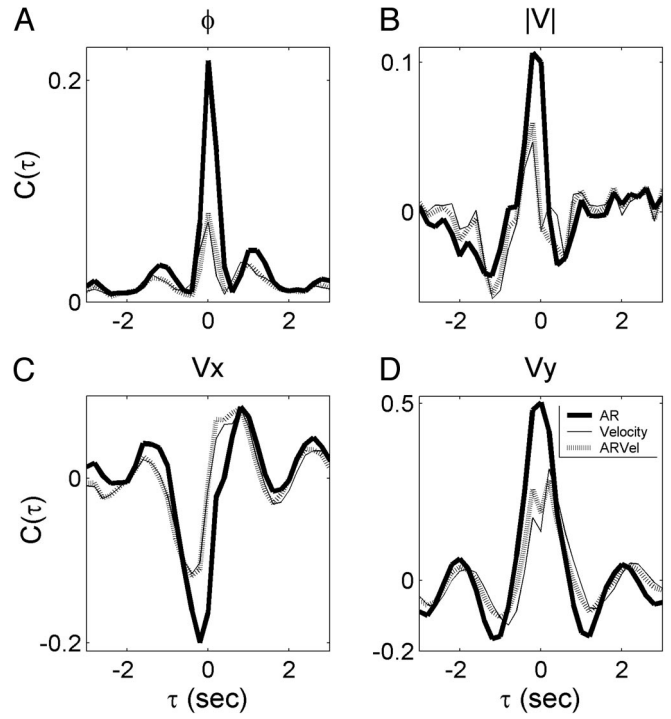

FIG. 7. Point process residual analysis. *A*: cross-correlation function $C(\tau)$ between the hand-movement direction and the residuals from the autoregressive spiking history (AR, thick curve), the velocity (thin curve), and the autoregressive spiking history plus velocity (ARVel, dashed curve) models applied to cell 75a. *B–D*: cross-correlations functions between the residuals and speed, and velocity in Cartesian coordinates (V_x and V_y). Correlations are significantly reduced for the velocity model in comparison to the autoregressive spiking history model. Nonetheless, there remains some structure in the point process residual that is related to the hand velocity but was not captured by the velocity model. Correlations for the velocity model were practically identical to the autoregressive spiking history plus velocity model, suggesting that the autoregressive component does not provide additional information about velocity (see text for details).

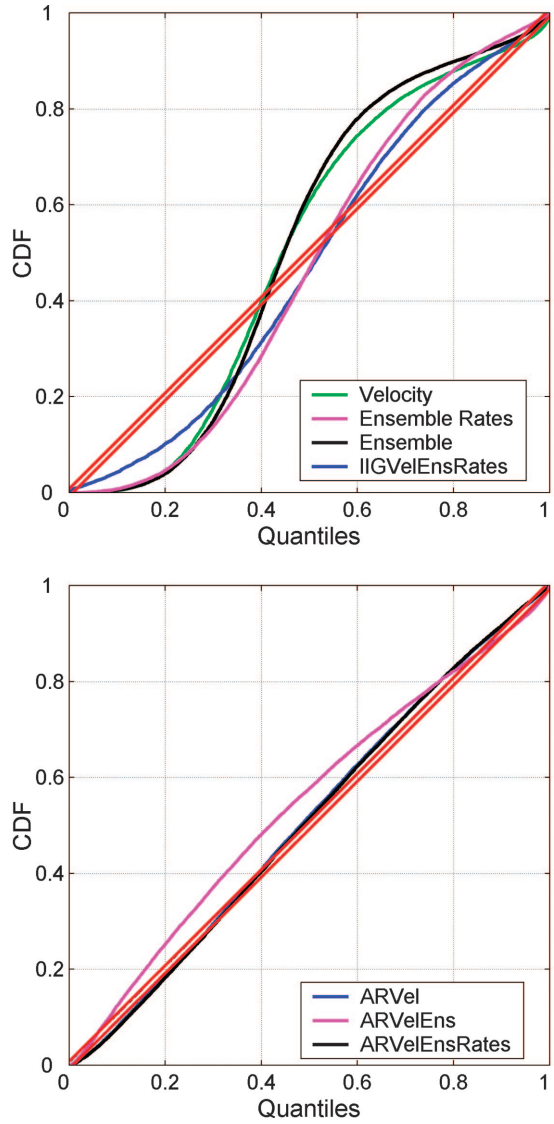

FIG. 8. Goodness-of-fit assessment (K-S plots) of alternative models. For comparison purposes, the models are shown in 2 groups. *Top*: the velocity, ensemble, ensemble rates, and the inhomogeneous inverse Gaussian plus velocity and ensemble rates models (IIGVelEnsRates) are compared. *Bottom*: the autoregressive spiking history plus velocity model (ARVel) is compared to 2 other models that add the ensemble (ARVelEns) or the ensemble rates component (ARVelEnsRates). K-S plots for the ARVel and the ARVelEnsRates partially overlap. (Cell 75a).

fining the conditional intensity function defines a point process model, the explanatory contributions of the covariates were assessed by constructing conditional intensity function models in terms of these covariates and by using a likelihood-based estimation approach to fit these models to data and to assess goodness-of-fit.

Analyses that measure the simultaneous effects of different covariates are crucial because the covariates modulate the neural spiking activity at the same time. Currently used analysis methods do not allow for the modeling of the simultaneous effects of spiking history, neural ensemble, and extrinsic covariates on the spike train treated as a point process. To evaluate the relation between spiking activity and covariates, the spike train data are frequently transformed to a rate function and the relation between the rate function and the covariate is then assessed using regression methods (e.g., Ashe and Georgopoulos 1994; Luczak et al. 2004). The relation between a neuron's spiking activity, spiking history, and concurrent neural activity are usually assessed using autocorrelation and pairwise cross-correlation analyses performed directly on the spike train (e.g., Hatsopoulos et al. 1998). The use of different methods to assess individually the importance of these covariates precludes an analysis of the neural point process in which the relative importance of all covariates is assessed and may also lead to a misleading estimate of the covariate effects. For example, spiking history effects can interfere with the accurate estimation of extrinsic covariate effects in spike triggered averages and reverse correlation methods (Aguera y Arcas and Fairhall 2003; Aguera y Arcas et al. 2003). Additionally, current analysis techniques also assess the contribution of ensemble covariates separately. For instance, pairwise cross-correlation analyses measure the statistical association of a single neuron to each member of the ensemble separately but not the association between the single neuron's activity and the entire observed ensemble.

The key to our likelihood approach is representing the conditional intensity function of a single neuron in terms of the covariates. In this way, the covariates are directly related to the probability that the neuron spikes. Although this formulation can be used generally to analyze the relation between covariates and neural spiking activity, it usually requires writing a new algorithm or function to carry out the maximum likelihood

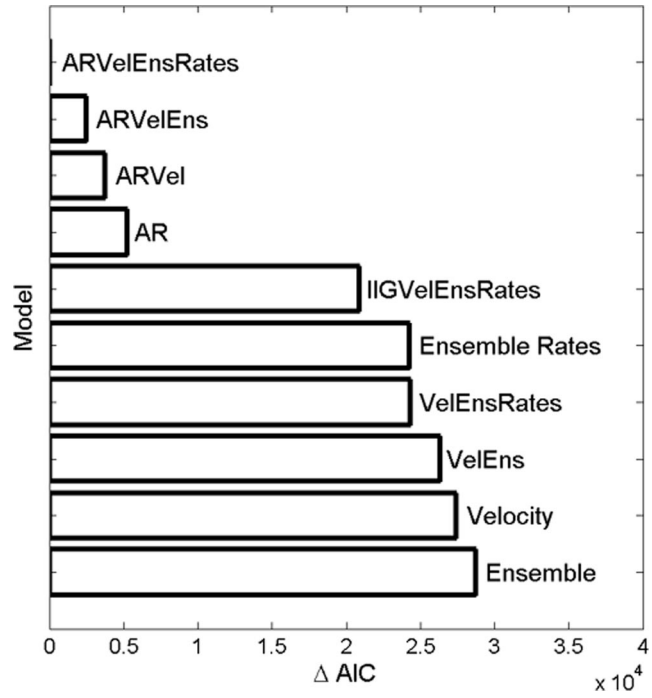

FIG. 9. Akaike's standard information criterion (AIC) model comparison. For convenience, we plot the differences of the AICs, denoted by ΔAIC of all of the models with respect to the AIC of the best model. Following this criterion and convention, better models have smaller AIC differences. Model labels: autoregressive spiking history (AR), autoregressive spiking history plus velocity (ARVel), autoregressive spiking history plus velocity and ensemble (ARVelEns), autoregressive spiking history plus velocity and ensemble rates (ARVelEnsRates), inhomogeneous inverse Gaussian plus velocity and ensemble rates models (IIGVelEnsRates), velocity plus ensemble (VelEns), and velocity plus ensemble rates (VelEnsRates). See METHODS section for model details. (Cell 75a).

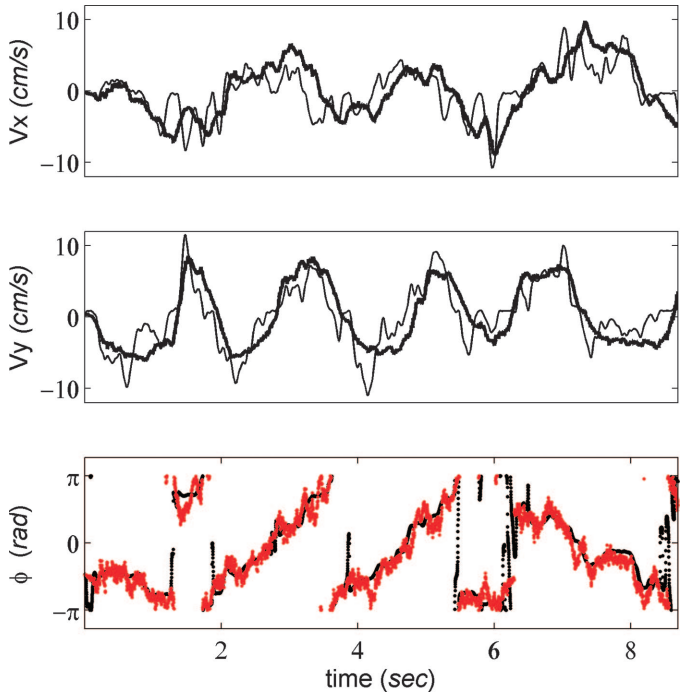

FIG. 10. Neural decoding of (x, y) velocities and movement direction by the point process filter. Estimated velocities (thick curve) and direction (red dots), together with true velocities and direction, are shown for a single decoded test trial. Time was discretized at a 1-ms resolution. At every millisecond, an estimate of the velocity parameters was obtained based on the state of the cells in the ensemble. All 20 recorded cells were used in the decoding. Conditional intensity function for each cell was given by the velocity model. Original decoding was done in polar coordinates. From a total of 130 trials, 120 trials were used for model fitting and 10 test trials for neural decoding. See Table 1 for summary statistics over the entire ensemble of test trials.

estimation with each new model formulation. To make our point process likelihood approach more broadly applicable, we showed that by representing the logarithm of the conditional intensity function as a linear combination of general functions of the covariates, the conditional intensity function model could be efficiently fit to neural spike train data using the GLM framework under a Poisson distribution and a log link function (*Eqs. 3*, *A6*, and *A8*). We also showed that, equivalently, if the point process is represented as a conditionally dependent Bernoulli process (*Eqs. 2* and *A2*) and the probability of the events is modeled by a logistic function, then the point process model could also be fit using the GLM framework under a Bernoulli distribution and a logistic link function (*Eqs. 5* and *A10*). Regarding the time discretization, we note that the 1-ms discretization interval we chose should not be interpreted as the absolute discretization interval for all applications of the point process–GLM framework. How fine the discretization interval should be will depend on the particular problem under study.

Our use of the Poisson distribution to fit point process models does not mean that we are assuming that our original data are Poisson. We are rather exploiting for computational purposes the fact that all point process likelihoods that admit a conditional intensity function, including those that are history dependent, have the same mathematical form given by *Eq. A6*. Our analysis makes explicit the relation between likelihood methods for point processes, conditional Bernoulli processes, and GLM model fitting with Poisson or Bernoulli distributions. This relation provides a justification for using this GLM framework to analyze spike trains as point process data.

The point process–GLM framework for analyzing neural spike train data has several important advantages. First, this framework allows us to formulate complex models to relate the spiking activity to covariates. The fact that *Eq. 4* is written in terms of general functions of the covariates provides the framework with a very large class of possible models. Second, the GLM framework is part of several standard mathematical and statistical packages (e.g., Splus, Matlab, SPSS), so that the approach is readily accessible to experimentalists analyzing their data. Even though likelihood-based methods are highly desirable because of several optimality properties, the biggest impediment to their widespread use is the lack of readily available software in which a flexible class of neural spike train models can be easily applied to data. The point process–GLM framework offers a practical, broadly applicable solution to the computational problem of fitting potentially complex point process models for neural spike trains by maximum likelihood.

Third, the point process–GLM framework makes it possible to apply a set of goodness-of-fit tools for point processes not available in the GLM. These are the point process residuals analysis, goodness-of-fit tests based on the time-rescaling theorem, and decoding from point process observations. Fourth, the model selection and goodness-of-fit methods available in GLM are extended to spike train data. Thus we have a set of complementary methods to assess the extent to which proposed models explain the structure in neural spike trains. Although we used the point process–GLM framework to carry out most of the analysis, we also illustrated with the conditional intensity function based on the inhomogeneous inverse Gaussian model how non-GLM point process likelihood models may be used to analyze neural spike train data. Finally, by analogy with the way in which linear regression methods are used to analyze the relation between a continuous dependent variable and a set of candidate explanatory variables, the ready availability of software to implement GLMs also makes it possible for neurophysiologists to quickly assess the relevance of a wide range of covariates before proceeding to construct more specific models that may require non-GLM algorithms to carry out the model fitting.

A key objective of the proposed framework is to provide tools for the assessment of the relative importance of the covariates on neural spiking activity. We showed how this objective is accomplished by analyzing the goodness-of-fit of

TABLE 1. *Mean coverage probabilities, mean, and median errors for the velocity, and the autoregressive spiking history plus velocity models*

	Velocity	(AR) History + Velocity
$(\|V\|, \phi)$: Mean cov. prob.	0.30 ± 0.10	0.24 ± 0.08
$\|V\|$: Mean cov. prob.	0.32 ± 0.10	0.27 ± 0.08
ϕ: Mean cov. prob.	0.94 ± 0.02	0.90 ± 0.04
(V_x, V_y): Mean error	3.60 ± 0.3	3.60 ± 0.4
(V_x, V_y): Median error	3.43	3.38
(V_x, V_y): Error (95% CI)	(0, 7.8]	(0, 7.8]
ϕ: Mean error	$0.15\pi \pm 0.1$	$0.14\pi \pm 0.1$
ϕ: Median error	0.11π	0.11π
ϕ: Error (95% CI)	$(0, 0.65\pi]$	$(0, 0.63\pi]$

CI denotes the confidence interval. Velocity is given in cm/s.

each model we proposed. Because no method provides a complete assessment of goodness-of-fit, we used the standard statistical approach of using multiple complementary measures that evaluate different aspects of the model's agreement with the data. We applied 4 complementary types of goodness-of-fit techniques in our example data set: Kolmogorov–Smirnov tests based on the time-rescaling theorem, model selection by AIC, point process residual analysis, and neural spike train decoding based on point process observations.

First, the K-S analyses performed on the example data problems were useful in providing a sense of how close our models were to capturing the stochastic structure in the example data sets. The fact that some of our models captured most of the statistical structure in the spiking activity (Figs. 4 and 8) suggests that developing a parsimonious statistical model of MI activity is a realistic goal. The K-S plots also highlighted the importance of the spiking history. The fact that the coefficients had a structure that seemed to reflect mostly refractory and recovery periods together with rebound excitation in short time scales (effectively shorter than 60 ms) suggests that the autoregressive component successfully captured important history effects.

Second, the AIC analysis provided additional information for model comparison when the K-S plots did not distinguish between different models. By finely ranking the different models, the AIC analysis allowed for the assessment of the distinct effects and relative importance of the 3 types of covariates. Our choice of the AIC for model comparison was motivated by the fact that AIC was derived to be an estimate of the expected relative Kullback–Leibler distance between the distributions generated by the model and the distribution of the underlying stochastic process generating the data. The "true" underlying model does not need to be in the set of tested models for the AIC analysis to suggest the most appropriate model, a situation that we believe is more often the rule than the exception in biology. Because the AIC is penalized with increasing numbers of model parameters, its use is more appropriate than the use of the data likelihood itself in preventing overfitting of the data by the model. Additionally, an equivalence between AIC and cross-validation in model comparison problems has been established (Stone 1977), and model-based penalty methods (e.g., AIC) outperform cross-validation in important aspects when assessing prediction performance (Efron 2004). We also computed the Bayesian information criterion (BIC) (Schwarz 1978), an alternate, more conservative criterion function for each of our models and found that it yielded the same model rankings as did the AIC. When comparing complex models, another protection against overfitting comes from having large quantities of data as compared to the number of model parameters. In our example the $>10^6$ data observations far outnumbered the maximum 200 or so parameters in our most complex models. Apparent inconsistencies between AIC and K-S analyses could occur in cases where, for example, the conditional intensity model is more likely to produce the observed spike train as a whole, but is less accurate in describing a specific aspect of the data structure, such as the regime of small ISIs. This might have been the case when comparing the AIC and K-S plots results for the autoregressive spiking history plus velocity and ensemble model and a simpler autoregressive spiking history plus velocity model (Figs. 8 and 9).

Third, we illustrated how the point process residual analysis can be used to assess the contribution of an extrinsic covariate to a single neuron's spiking activity. In the illustration example, the residual analysis showed that the introduction of the velocity covariate captured a significant amount of the statistical structure related to hand velocity available in the spiking activity of a single neuron. Yet, for the particular cell chosen in this study, the analysis was also able to reveal that there still was a significant amount of structure in the residuals that was correlated to hand velocity but that was not captured by this specific form of the velocity model. Cross-correlation analysis of the point process residuals and extrinsic covariates is thus an important tool for assessing whether a particular model has captured well the effects of the covariate on the spiking activity. The ideal model should produce a residual with no significant correlations to the modeled covariate. It should also be noted that, unlike the decoding analysis, the point process residual analysis is not dependent on the properties of a decoding algorithm.

Fourth, complementing the above 3 goodness-of-fit analysis tools, the spike train decoding allowed for the goodness-of-fit assessment at the neural ensemble level. In conjunction with understanding what makes a neuron spike, we are also interested in assessing how well a model captures the representation of an extrinsic covariate at the ensemble level. At present, decoding is the only technique we have for assessing goodness-of-fit at this level. The key elements for assessing goodness-of-fit in the decoding analysis were the predicted signal and its confidence intervals and coverage probability, and especially the estimation error and its confidence intervals. The confidence intervals and the coverage probability based on the estimated posterior covariance matrix provided an estimate of the amount of uncertainty in the decoding algorithm, whereas the decoding error and its distribution provided a measure of the algorithm's actual performance. For this reason, the mean coverage probability should be interpreted in conjunction with the mean decoding error. As suggested by the narrow distribution of the decoding error and approximately 0.95 mean coverage probability (Table 1), hand movement direction was remarkably well decoded. Velocity estimates in Cartesian coordinates captured reasonably well slow fluctuations in the measured hand velocity. We also illustrated how to assess the contribution of the autoregressive spiking history component to neural decoding. The autoregressive spiking history plus velocity and the velocity models performed similarly well. This preliminary result suggests that short time dynamics captured by the autoregressive component did not play a crucial role in decoding hand velocity or movement direction in this data set. Given that the models used in the decoding analysis did not include the ensemble or ensemble rates covariate, the ensemble decoding assumed independent encoder cells. Nonetheless, the framework also allows the assessment of the contribution of interaction patterns in the neural ensemble to decoding. That could be easily achieved by extending the conditional intensity models to include the ensemble covariates. This analysis is beyond the scope of this paper and will be addressed elsewhere.

In summary, the above 4 complementary goodness-of-fit and model selection analyses are an essential step for achieving our primary objective of assessing the effects and relative importance of the modeled covariates. The proposed point process

framework provides a starting point for building and analyzing complex models of spiking activity including the 3 types of covariates discussed in this paper. In particular, the point process–GLM framework provides a systematic approach to neural spike train data analysis analogous to that for continuous-valued variables under the linear Gaussian regression framework.

We foresee several potential improvements and extensions of this framework. Although we were able to fit models with hundreds of parameters, much larger models will require the development of efficient algorithms for both GLM and non-GLM computations. Also, to analyze the relation between the simultaneous firing of an entire ensemble relative to its spiking history and a set of extrinsic covariates, we can extend the framework by using multivariate point process likelihoods (Chornoboy et al. 1988). A multivariate likelihood model will facilitate the study of the independence, redundancy, and synergy in the ensemble representation. Finally, multielectrode devices (Csicsvari et al. 2003; Donoghue 2002; Nicolelis et al. 2003) now make possible the simultaneous recordings of multiple single cells from many different brain areas. We foresee the proposed framework as a valuable tool for investigating how interacting brain regions represent, compute, and process information. We are currently applying this framework to the analysis of parietal and MI spiking activity in monkeys performing visuomotor tracking tasks and hippocampus activity in rats performing a range of learning tasks.

APPENDIX

Continuous and discrete time point process likelihood function

The likelihood of a neural spike train, like that of any statistical model, is defined by finding the joint probability density of the data. We show below that the joint probability of any point process is easy to derive from the conditional intensity function. We show that the point process likelihood function in *Eqs. 2* and *3* gives a discrete time approximation of the likelihood function for a continuous time point process (*Eq. A6* below).

Let $0 < u_1 < u_2, \ldots, u_J < T$ be a set of neural spike train measurements. Using the discrete time representation given in the METHODS section, define the events

$$A_k = \{\text{spike in } (t_{k-1}, t_k] \mid H_k\}$$
$$E_k = \{A_k\}^{\Delta N_k}\{A_k^c\}^{1-\Delta N_k}$$
$$H_k = \left\{\bigcap_{j=1}^{k-1} E_j\right\} \tag{A1}$$

for $k = 1, \ldots, K$ and where A_k^c is the complement of A_k. For simplicity, H_k includes only the intrinsic history of the process. It can be easily extended to incorporate neural ensemble activity and other extrinsic covariates. By construction of the partition of the interval (0, T], introduced in the METHODS section, we must have $u_j \in (t_{k_j-1}, t_{k_j}]$, $j = 1, \ldots, J$, for a subset of the intervals satisfying $k_1 < k_2 < \cdots < k_J$. The remaining $K - J$ intervals have no spikes.

The probability of exactly J events occurring within the intervals $(t_{k_j-1}, t_{k_j}]_{j=1}^{J}$ in (0, T], may then be computed as

$$P(N_{1:K}) = P(u_j \in (t_{k_j-1}, t_{k_j}], j = 1, \ldots, J, \cap N(T) = J)$$
$$= \prod_{k=1}^{K} P(A_k)^{\Delta N_k} P(A_k^c)^{1-\Delta N_k} \tag{A2}$$

by the definition of A_k and E_k in *Eq. A1*.

The spike train thus forms a sequence of conditionally independent Bernoulli trials, with the probability of a spike in the k^{th} time interval given by $P(A_k)$. In any interval $(t_{k-1}, t_k]$ we have

$$P(A_k) = \lambda(t_k \mid H_k)\Delta + o(\Delta)$$
$$P(A_k^c) = 1 - \lambda(t_k \mid H_k)\Delta + o(\Delta) \tag{A3}$$

Substituting *Eq. A3* into *Eq. A2* yields

$$P(N_{1:K}) = \prod_{k=1}^{K} [\lambda(t_k \mid H_k)\Delta]^{\Delta N_k}[1 - \lambda(t_k \mid H_k)\Delta]^{1-\Delta N_k} + o(\Delta^J) \tag{A4}$$

which is *Eq. 2*. For small Δ, $[1 - \lambda(t_k)\Delta] \approx \exp\{-\lambda(t_k)\Delta\}$ and $\log [\lambda(t_k)\Delta[1 - \lambda(t_k)\Delta]^{-1}] \approx \log (\lambda(t_k)\Delta)$, therefore we obtain

$$\begin{aligned} P(N_{1:K}) &= \prod_{k=1}^{K} [\lambda(t_k \mid H_k)\Delta]^{\Delta N_k}[1 - \lambda(t_k \mid H_k)\Delta]^{-\Delta N_k} \prod_{k=1}^{K} [1 - \lambda(t_k \mid H_k)\Delta] + o(\Delta^J) \\ &= \prod_{k=1}^{K} \left[\frac{\lambda(t_k \mid H_k)\Delta}{1 - \lambda(t_k \mid H_k)\Delta}\right]^{\Delta N_k} \prod_{k=1}^{K} \exp\{-\lambda(t_k \mid H_k)\Delta\} + o(\Delta^J) \\ &= \exp\left\{\sum_{k=1}^{K} \log [\lambda(t_k \mid H_k)\Delta]\Delta N_k - \sum_{k=1}^{K} \lambda(t_k \mid H_k)\Delta\right\} + o(\Delta^J) \end{aligned} \tag{A5}$$

The probability density of these J exact spikes in (0, T], given by $p(N_{0:T}) = \lim_{\Delta\to 0} P(N_{1:K})/\Delta^J$, is then obtained as

$$\begin{aligned} P(N_{0:T}) &= \lim_{\Delta\to 0} \frac{\exp\left\{\sum_{k=1}^{K} \log [\lambda(t_k \mid H_k)\Delta]\Delta N_k - \sum_{k=1}^{K} \lambda(t_k \mid H_k)\Delta\right\} + o(\Delta^J)}{\Delta^J} \\ &= \lim_{\Delta\to 0} \frac{\exp\left\{\sum_{k=1}^{K} \log \lambda(t_k \mid H_k)\Delta N_k - \sum_{k=1}^{K} \lambda(t_k \mid H_k)\Delta\right\}\Delta^J + o(\Delta^J)}{\Delta^J} \\ &= \exp\left\{\int_0^T \log \lambda(t \mid H(t))\mathrm{d}N(t) - \int_0^T \lambda(t \mid H(t))\mathrm{d}t\right\} \end{aligned} \tag{A6}$$

which is the joint probability density of the point process spike train in continuous time (Brown et al. 2003; Daley and Vere-Jones 2003). Note that we could have derived the likelihood for the continuous time point process (and therefore also *Eq. 3*) by a generalization of the continuous time Poisson process (Daley and Vere-Jones 2003), without resorting to representing the neural point process as a conditional Bernoulli process. We formulated the spike train joint probability in terms of *Eq. A2* only to show (see below) the equivalence between Poisson and Bernoulli–GLMs when Δ is sufficiently small.

The Poisson and Bernoulli–GLMs

We briefly define a generalized linear model and show that for small enough Δ, the Bernoulli and Poisson–GLMs are equivalent in the modeling of spiking train data.

Two main aspects characterize a generalized linear model of a random variable y (McCullagh and Nelder 1989). First, the modeled random variable y has a distribution in the exponential family. Among several members of this family are the Gaussian, the Poisson, and the Bernoulli distribution. The exponential family has the general form

$$f(y \mid \Theta, \phi) = \exp\{[y\Theta - b(\Theta)]/a(\phi) + c(y, \phi)\} \quad (A7)$$

where $a(\cdot)$, $b(\cdot)$, and $c(\cdot)$ are some specific functions. If ϕ is known, this is an exponential-family model with canonical parameter Θ.

For the particular case of the Poisson distribution, $\Theta = \log\lambda$, $b(\Theta) = e^{\Theta}$, $a(\phi) = \phi = 1$, $c(y, \phi) = -\log(y!)$. The location and scale parameters are λ and ϕ, respectively. Thus the distribution in *Eq. A7* can be expressed as $f(y \mid \lambda) = \exp\{y \log\lambda - \lambda - \log(y!)\} = \lambda^y e^{-\lambda}/y!$. Note that the canonical parameter Θ has, in the Poisson case, a natural representation in terms of the logarithm of the parameter λ. The joint probability distribution for an independently distributed data set $\mathbf{y} = \{y_k\}_{k=1}^K$ becomes

$$f(\mathbf{y} \mid \Theta, \phi) \propto \exp\left\{\sum_{k=1}^{K} y_k \log\lambda_k - \sum_{k=1}^{K} \lambda_k\right\} \quad (A8)$$

If the rate function λ_k of this Poisson process is generalized by the conditional intensity function (*Eq. 1*); and $y_k = \Delta N_k$, $\Delta = 1$, then *Eq. A8* has the same form as the typical general likelihood function for any discrete time point process (*Eqs. 3* and *A5*).

For the Bernoulli case, we let $y = \{0, 1\}$ with the probability of success denoted by P, and set $\Theta = \log([1 - P]^{-1}P)$, $b(\Theta) = \log(1 + e^{\Theta})$, $a(\phi) = \phi$, $c(\cdot) = 1$, and $\phi = 1$. Thus, for single realizations we have $p(y \mid P) = P^y(1 - P)^{1-y}$. Given an independently distributed data set $\mathbf{y} = \{y_k\}_{k=1}^K$, the likelihood function under the Bernoulli distribution becomes

$$f(\mathbf{y} \mid \Theta, \phi) = \prod_{k=1}^{K} [P_k]^{y_k}[1 - P_k]^{1-y_k} \quad (A9)$$

By letting $[P_k]^{y_k} = P(A_k)^{\Delta N_k}$ we obtain *Eq. A2*.

Second, the defining feature of a generalized linear model follows. The canonical parameter of the exponential family is expressed in a linear form by a transformation given by a monotonic differentiable function. In the Poisson case, if the canonical parameter is modeled as a linear combination of general functions of covariates $\mathbf{v}$ of interest (that is, $\Theta = \log\lambda(\theta, \mathbf{v}) = \Sigma_{i=1}^{q} \theta_i g_i(v_i)$ as in *Eq. 4* or equivalently as $\lambda(\theta, \mathbf{v}) = \exp\{\Sigma_{i=1}^{q} \theta_i g_i(v_i)\}$ as in *Eqs. 5–10*), then $f(y \mid \Theta, \phi) \propto \exp\{y \log\lambda(\theta, \mathbf{v}) - \lambda(\theta, \mathbf{v})\}$ gives the distribution for a GLM under a Poisson distribution and a log link function. In the Bernoulli case, if $\Theta = \log([1 - P(A_k \mid \theta, H_k)]^{-1}P(A_k \mid \theta, H_k))$ is modeled as linear combination of general functions of the covariates, then *Eqs. 2* and *A2* give the likelihood function for a GLM under a Bernoulli distribution and a *logistic* link function.

Finally, we establish the relation between the Poisson and Bernoulli–GLMs in the context of neural point process models. After making explicit the parametric model of the conditional intensity function, we have the probability of a spike event in the time interval $(t_{k-1}, t_k]$ given by $P(A_k \mid \theta, H_k) = \lambda(t_k \mid \theta, H_k)\Delta + o(\Delta)$. For small Δ

$$\log\left[\frac{P(A_k \mid \theta, H_k)}{1 - P(A_k \mid \theta, H_k)}\right] \approx \log[\lambda(t_k \mid \theta, H_k)\Delta] \quad (A10)$$

Therefore, for small enough Δ, the Bernoulli and Poisson–GLMs are equivalent when applied to the modeling of spiking train data.

The IIG model for the ISI probability density

For a particular cell, let $t_e = t_k - u_{N_{k-1}}$ denote the time elapsed since the last spike $u_{N_{k-1}}$. The inhomogeneous inverse Gaussian ISI probability density function conditioned on the motor covariate and neural ensemble activity is defined as

$$p(t_e \mid G) = \frac{s(t_k \mid \cdot)}{\left[2\pi\left[\int_{u_{N_{k-1}}}^{t_k} s(t \mid \cdot)\mathrm{d}t\right]^3\right]^{1/2}} \exp\left\{-\frac{1}{2}\frac{\left(\int_{u_{N_{k-1}}}^{t_k} s(t \mid \cdot)\mathrm{d}t - \psi\right)^2}{\psi^2 \int_{u_{N_{k-1}}}^{t_k} s(t \mid \cdot)\mathrm{d}t}\right\} \quad (A11)$$

where $G = \{u_{N_{k-1}}, \mathbf{x}_{k+\tau}, N_{1:k}^{1:C}, \theta\}$, $\psi > 0$ is the location parameter and $s(t_k \mid \cdot) > 0$ is the scaling parameter at time t_k, conditioned on the extrinsic and ensemble rates covariates as given in *Eq. 13*.

GRANTS

This work was supported in part by Burroughs Wellcome Fund, Keck Foundation, National Institute of Neurological Disorders and Stroke Neural Prosthesis Program, National Institutes of Health Grants NS-25074, MH-59733, MH-61637, MH-65018, and DA-015644, National Science Foundation Grant 0081548, and grants from Defense Advanced Research Projects Agency and Office of Naval Research.

DISCLOSURE

J. P. Donoghue is a cofounder and shareholder in Cybernetics, Inc., a neurotechnology company that is developing neural prosthetic devices.

REFERENCES

Aguera y Arcas B, and Fairhall AL. What causes a neuron to spike? *Neural Comput* 15: 1789–1807, 2003.

Aguera y Arcas B, Fairhall AL, and Bialek W. Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput* 15: 1715–1749, 2003.

Akaike H. Information theory as an extension of the maximum likelihood principle. In: *Second International Symposium on Information Theory*, edited by Petrov BN and Csaki F. Budapest: Akademiai Kiado, 1973, p. 267–281.

Andersen PK, Borgan O, Gill RD, and Keiding N. *Statistical Models Based on Counting Processes.* New York: Springer-Verlag, 1992.

Ashe J and Georgopoulos AP. Movement parameters and neural activity in motor cortex and area-5. *Cereb Cortex* 4: 590–600, 1994.

Barbieri R, Quirk MC, Frank LM, Wilson MA, and Brown EN. Construction and analysis on non-Poisson stimulus-response models of neural spiking activity. *J Neurosci Methods* 105: 25–37, 2001.

Berman M and Turner TR. Approximating point process likelihoods with GLIM. *Appl Stat* 41: 31–38, 1992.

Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol Cybern* 59: 189–200, 1988.

Brown EN, Barbieri R, Eden UT, and Frank LM. Likelihood methods for neural data analysis. In: *Computational Neuroscience: A Comprehensive Approach*, edited by Feng J. London: CRC, 2003.

Brown EN, Barbieri R, Ventura V, Kass RE, and Frank LM. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput* 14: 325–2346, 2002.

Burnham KP and Anderson DR. *Model Selection and Multimodel Inference: A Practical Information Theoretic Approach* (2nd ed.). New York: Springer-Verlag, 2002.

Chornoboy ES, Schramm LP, and Karr AF. Maximum likelihood identification of neural point process systems. *Biol Cybern* 59: 265–275, 1988.

Cover TM and Thomas JA. *Elements of Information Theory.* New York: Wiley, 1991.

Csicsvari J, Henze DA, Jamieson B, Harris KD, Sirota A, Bartho P, Wise KD, and Buzsaki G. Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J Neurophysiol* 90: 1314–1323, 2003.

Daley D and Vere-Jones D. *An Introduction to the Theory of Point Processes.* New York: Springer-Verlag, 2003.

Donoghue JP. Connecting cortex to machines: recent advances in brain interfaces. *Nat Neurosci* 5: 1085–1088, 2002.

Donoghue JP, Sanes JN, Hatsopoulos NG, and Gaal G. Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol* 79: 159–173, 1998.

Eden UT, Frank LM, Barbieri R, Solo V, and Brown EN. Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Comput* 16: 971–998, 2004.

Efron B. The estimation of prediction error: Covariance penalties and cross-validation. *J Am Stat Assoc* 99: 619–632, 2004.

Frank LM, Eden UT, Solo V, Wilson MA, and Brown EN. Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: an adaptive filtering approach. *J Neurosci* 22: 3817–3830, 2002.

Fu QG, Flament D, Coltz JD, and Ebner TJ. Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J Neurophysiol* 73: 836–854, 1995.

Grammont F and Riehle A. Precise spike synchronization in monkey motor cortex involved in preparation for movement. *Exp Brain Res* 128: 118–122, 1999.

Grammont F and Riehle A. Spike synchronization and firing rate in a population of motor cortical neurons in relation to movement direction and reaction time. *Biol Cybern* 88: 360–373, 2003.

Hatsopoulos NG, Ojakangas CL, Paninski L, and Donoghue JP. Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proc Natl Acad Sci USA* 95: 15706–15711, 1998.

Hatsopoulos NG, Paninski L, and Donoghue JP. Sequential movement representations based on correlated neuronal activity. *Exp Brain Res* 149: 478–486, 2003.

Hille B. *Ion Channels of Excitable Membrane* (3rd ed.). Sunderland, MA: Sinauer Associates, 2001.

Jackson A, Gee VJ, Baker SN, and Lemon RN. Synchrony between neurons with similar muscle fields in monkey motor cortex. *Neuron* 38: 115–125, 2003.

Johnson A and Kotz S. *Distributions in Statistics: Continuous Univariate Distributions.* New York: Wiley, 1970.

Kass RE and Ventura V. A spike-train probability model. *Neural Comput* 13: 1713–1720, 2001.

Keat J, Reinagel P, Reid RC, and Meister M. Predicting every spike: a model for the responses of visual neurons. *Neuron* 30: 803–817, 2001.

Kitagawa G and Gersh W. *Smoothness Priors Analysis of Time Series.* New York: Springer-Verlag, 1996.

Knierim JJ and Vanessen DC. Neuronal responses to static texture patterns in area-V1 of the alert macaque monkey. *J Neurophysiol* 67: 961–980, 1992.

Luczak A, Hackett TA, Kajikawa Y, and Laubach M. Multivariate receptive field mapping in marmoset auditory cortex. *J Neurosci Methods* 136: 77–85, 2004.

Maynard EM, Hatsopoulos NG, Ojakangas CL, Acuna BD, Sanes JN, Normann RA, and Donoghue JP. Neuronal interactions improve cortical population coding of movement direction. *J Neurosci* 19: 8083–8093, 1999.

McCullagh P and Nelder JA. *Generalized Linear Models* (2nd ed.). Boca Raton, FL: Chapman & Hall/CRC, 1989.

Mehta MR, Barnes CA, and McNaughton BL. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci USA* 94: 8918–8921, 1997.

Mehta MR, Quirk MC, and Wilson MA. Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25: 707–715, 2000.

Mendel JM. *Lessons in Estimation Theory for Signal Processing, Communications, and Control.* Englewood Cliffs, NJ: Prentice Hall, 1995.

Moran DW and Schwartz AB. Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82: 2676–2692, 1999.

Nicolelis MA, Dimitrov D, Carmena JM, Crist R, Lehew G, Kralik JD, and Wise SP. Chronic, multisite, multielectrode recordings in macaque monkeys. *Proc Natl Acad Sci USA* 100: 11041–11046, 2003.

Ogata Y. On Lewis' simulation method for point processes. *IEEE Trans Inform Theory* IT-27: 23–31, 1981.

Ogata Y. Statistical models for earthquake occurrences and residual analysis for point processes. *J Am Stat Assoc* 83: 9–27, 1988.

O'Keefe J and Dostrovsky J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34: 171–175, 1971.

O'Keefe J and Recce ML. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3: 317–330, 1993.

Paninski L, Fellows MR, Hatsopoulos NG, and Donoghue JP. Spatiotemporal tuning of motor neurons for hand position and velocity. *J Neurophysiol* 91: 515–532, 2004.

Papangelou F. Integrability of expected increments of point processes and a related change of scale. *Trans Am Math Soc* 165: 483–506, 1972.

Pawitan Y. *In All Likelihood: Statistical Modeling and Inference Using Likelihood.* London: Oxford Univ. Press, 2001.

Sanes JN and Truccolo W. Motor "binding": do functional assemblies in primary motor cortex have a role? *Neuron* 38: 3–5, 2003.

Santner TJ and Duffy DE. The statistical analysis of discrete data. New York: Springer-Verlag, 1989.

Schwarz G. Estimating the dimension of a model. *Ann Stat* 6: 461–464, 1978.

Scott SH. The role of primary motor cortex in goal-direction movements: insights from neurophysiological studies on non-human primates. *Curr Opin Neurobiol* 13: 671–677, 2003.

Sillito AM, Grieve KL, Jones HE, Cudeiro J, and Davis J. Visual cortical mechanisms detecting focal orientation discontinuities. *Nature* 378: 492–496, 1995.

Stone M. An asymptotic equivalence of choice of model by cross-validation and Akaike's criterion. *J R Stat Soc B* 39: 44–47, 1977.

Tanner MA. *Tools for Statistical Inference: Methods for the Exploration of Posterior Distributions and Likelihood Functions.* New York: Springer-Verlag, 1996.

Vinje WE and Gallant JL. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287: 1273–1276, 2000.

Wilson HR. Simplified dynamics of human and mammalian neocortical neurons. *J Theor Biol* 200: 375–388, 1999.

Wilson MA and McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058, 1993.

nature neuroscience

Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes

Wilson Truccolo[1,2,5], Leigh R Hochberg[2–6] & John P Donoghue[1,2,4]

Coordinated spiking activity in neuronal ensembles, in local networks and across multiple cortical areas, is thought to provide the neural basis for cognition and adaptive behavior. Examining such collective dynamics at the level of single neuron spikes has remained, however, a considerable challenge. We found that the spiking history of small and randomly sampled ensembles (~20–200 neurons) could predict subsequent single neuron spiking with substantial accuracy in the sensorimotor cortex of humans and nonhuman behaving primates. Furthermore, spiking was better predicted by the ensemble's history than by the ensemble's instantaneous state (Ising models), emphasizing the role of temporal dynamics leading to spiking. Notably, spiking could be predicted not only by local ensemble spiking histories, but also by spiking histories in different cortical areas. These strong collective dynamics may provide a basis for understanding cognition and adaptive behavior at the level of coordinated spiking in cortical networks.

Single-neuron action potential (spiking) activity depends on intrinsic biophysical properties and the neuron's interactions in neuronal ensembles. In contrast with *ex vivo*/*in vitro* preparations, cortical pyramidal neurons in intact brain each commonly receive thousands of synaptic connections arising from a combination of short- and long-range axonal projections in highly recurrent networks[1–3]. Typically, a considerable fraction of these synaptic inputs is simultaneously active in behaving animals, resulting in 'high-conductance' membrane states[4]; that is, lower membrane input resistance and more depolarized membrane potentials. The large number of synaptic inputs and the associated high-conductance states contribute to the high stochasticity of spiking activity and the typically weak correlations observed among randomly sampled pairs of cortical neurons. Nevertheless, previous studies[5–11] have suggested that some of a single neuron's spiking activity might be explained by measured ongoing network states. We used cortical microelectrode array recordings in humans and monkeys to determine the predictability of single-neuron spiking on the basis of the recent (<100 ms) spiking history of small, randomly sampled neuronal ensembles from the same (intra) or from a different (inter), but connected, cortical area. In addition, we also compared the predictive power of ensemble spiking histories and instantaneous collective states. Substantial predictability in these small and randomly sampled ensembles would imply strong collective dynamics, with implications for both cortical processing and the experimental endeavor of studying coordinated spiking in large, distributed cortical networks.

We studied tens to hundreds of randomly and simultaneously sampled neurons in small (4 × 4 mm) patches in arm-related areas of primary motor (M1), parietal (5d) and ventral premotor (PMv) cortices while humans (M1) and monkeys (M1-PMv and M1-5d) performed sensorimotor tasks. Beyond their local connectivity, M1-PMv and M1-5d are known to be bidirectionally connected[12,13] (their coordination is thought be important for reaching and grasping[14,15]), allowing us to study not only local, but also inter-areal, ensemble–based prediction. Point process models[16–18] were fitted to express the spiking probability at any 1-ms time interval conditioned on the past 100 ms of the neuron's own (intrinsic) spiking history and the past 100 ms of the spiking history of the neuronal ensemble. On the basis of this conditional spiking probability, we predicted whether or not a target neuron would spike in any given 1-ms time bin. When examining the predictive power of the ensemble's simultaneous spiking state, we defined simultaneity at two levels of temporal resolution, 1 and 10 ms. These instantaneous collective states could result from common inputs or from synchronization patterns arising from the neuronal network's intrinsic dynamics. Pair-wise maximum entropy (Ising) models[7] were used to approximate the distributions of these instantaneous collective states. Detailed receiver operating characteristic (ROC) curve analyses[19] allowed us to quantify and compare the predictive power of intrinsic and ensemble histories, intra- and inter-areal ensemble activity, and spiking histories and instantaneous collective states.

RESULTS

We analyzed 12 neuronal datasets recorded from two human clinical trial participants with tetraplegia (hS1 and hS3) and four monkeys (mLA, mCL, mCO and mAB). These included M1 recordings taken while human participants performed two sessions of a 'neural cursor' center-out task (that is, a task where the participant used, via

[1]Department of Neuroscience, [2]Brown Institute for Brain Science and [3]Division of Engineering, Brown University, Providence, Rhode Island, USA. [4]Rehabilitation Research and Development Service, Department of Veterans Affairs, Providence, Rhode Island, USA. [5]Department of Neurology, Massachusetts General Hospital, Harvard Medical School, Boston, Massachusetts, USA. [6]Brigham and Women's Hospital, and Spaulding Rehabilitation Hospital, Harvard Medical School, Boston, Massachusetts, USA. Correspondence should be addressed to W.T. (wilson_truccolo@brown.edu).

Received 16 July; accepted 26 October; published online 6 December 2009; doi:10.1038/nn.2455

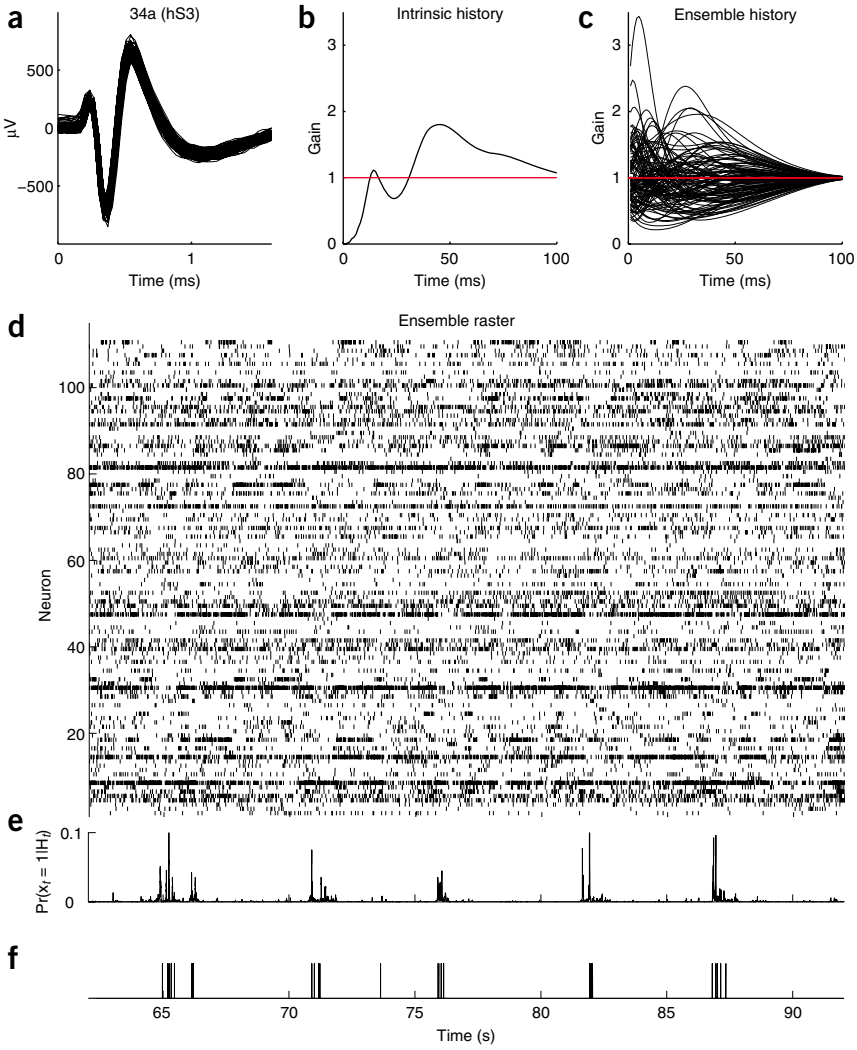

Figure 1 History point process models, intrinsic and ensemble history effects, and conditional spiking probabilities. Neuron 34a (hS3, session 2) was chosen as the example target neuron. (**a**) Waveforms corresponding to all sorted spikes for neuron 34a used in these analyses are shown. (**b**) Intrinsic spiking history. The curve represents the estimated temporal filter for the intrinsic history. Values below or above 1 correspond to a decrease or increase, respectively, in spiking probability contributed by a spike at a previous time specified in the horizontal coordinate. Refractory and recovery period effects after a spike, followed by an increase in spiking probability at longer time lags (40–100 ms), can be seen. This late intrinsic history effect might also reflect network dynamics. (**c**) Ensemble spiking history effects. Each curve represents the temporal filter corresponding to a particular input neuron to cell 34a. Many input neurons contributed biphasic effects: for example, an increase in spiking probability followed by a decrease, or vice-versa. All of the examined target neurons in our datasets showed qualitatively similar temporal filters. (**d**) Spike raster for all of the 110 neurons recorded in hS3 over a short, continuous time period. (**e**) Predicted spiking probabilities for the target neuron 34a were computed from the estimated intrinsic and ensemble temporal filters and the spike trains shown in **b**, **c** and **d**, respectively. (**f**) Observed spike train for neuron 34a in the same period.

a neural interface, his or her own recorded M1 spiking activity to move a computer cursor to targets radially positioned on the computer screen), M1 recordings from two monkeys (mLA and mCL), each performing two sessions of a task requiring planar point-to-point reaches toward targets randomly placed in the workspace, simultaneous M1 and PMv recordings from a monkey (mCO) performing two sessions of a task that required reach and grasp toward moving objects in a three-dimensional workspace, and simultaneous M1 and 5d recordings from a monkey (mAB) performing two sessions of a pursuit tracking task that required planar hand movements. The 12 datasets used in the analyses included 1,187 neuronal recordings. Minimum inter-electrode distance in the 10 × 10 microelectrode array was 400 μm. Maximum inter-electrode distance, including electrodes in two arrays, was ~2 cm (**Supplementary Fig. 1**).

To assess the predictive power of spiking histories, we first computed the probability that any given i^{th} neuron $x_{i,t}$ was going to spike at time t, conditioned on spiking histories $\mathcal{H}_t$ from the past 100 ms up to (but not including) time t. Without further constraints, direct estimation of conditional probability distributions for high-dimensional systems is typically an intractable problem, leading to combinatorial explosion and requiring amounts of data that grow exponentially with the number of neurons in the ensemble. Instead, we took advantage of the fact that this conditional probability can be computed as[16]

$$\Pr(x_{i,t} = 1 \mid \mathcal{H}_t) \approx \lambda_i(t \mid \mathcal{H}_t)\Delta \quad (1)$$

where $\lambda_i(t|\mathcal{H}_t)$ is the instantaneous spiking rate (conditional intensity function) of the neuron and $\Delta = 1$ ms in our discrete time representation, and used a simplified model to capture the relationship between the instantaneous rate and spiking histories

$$\log(\lambda_i(t \mid \mathcal{H}_t)\Delta) = \mu_i + K_{1,i} \cdot x_i + \sum_{j \neq i} K_{2,i,j} \cdot x_j \quad (2)$$

The term μ_i relates to a background level of spiking activity, x_i is the spiking history in the specified time interval for the i^{th} neuron, $i = 1, 2, \ldots, n$ recorded neurons, and $K_{1,i}$ and $K_{2,i,j}$ denote temporal filters related to intrinsic and ensemble history effects, respectively. These temporal filters were approximated via basis functions[9,16,18] (see Online Methods). Once an instantaneous spiking rate model was fitted, the estimated probability of a spike at any given time bin, conditioned on intrinsic and ensemble spiking histories, was easily computed using equation (1) (the spike prediction approach is shown using cell 34a (hS3) as the example target neuron; **Fig. 1**).

Predictive power of ensemble spiking histories in M1

On the basis of the estimated conditional spiking probabilities, we used a standard tool, the ROC curve analysis[19], to assess the predictive power of spiking history models (see Online Methods). Spike prediction in 1-ms time bins based on spiking histories was substantial. Intra-areal ensemble history–based predictions in M1 resulted in high true-positive rates while maintaining low numbers of false positives.

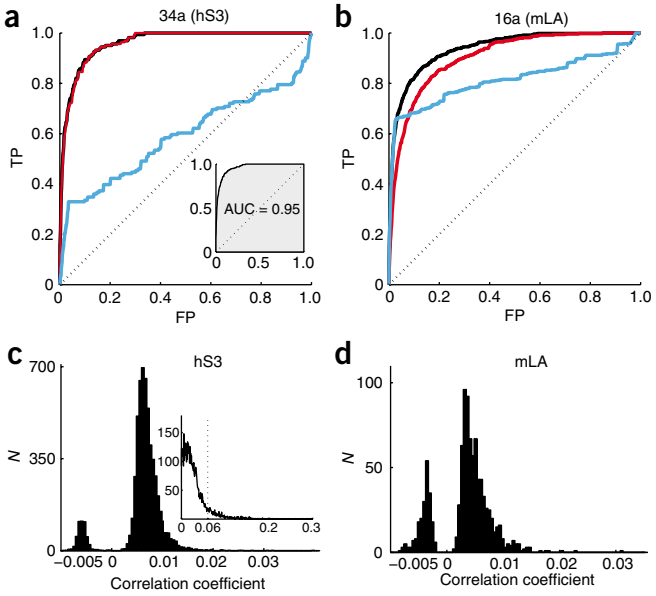

Figure 2 Prediction of single-neuron spiking and weak pair-wise correlations. (**a**) ROC curves for neuron 34a (human participant hS3, $n = 110$ neurons, session 2, 240,000 samples). FP and TP denote false- and true-positive prediction rates, respectively. The diagonal line corresponds to the expected chance prediction. The black, red and blue ROC curves correspond to the prediction based on full history models, only the ensemble histories, or only the neuron's own spiking history, respectively. The inset shows the AUC corresponding to the ROC curve for the ensemble history model. (**b**) ROC curves for neuron 16a (monkey mLA, $n = 45$, session 2, 1,230,857 samples). 95% confidence intervals for the AUC chance level resulted in 0.51 ± 0.004 and 0.51 ± 0.017 for target neurons 16a and 34a, respectively. These narrow confidence intervals (data not shown) were typical for the recorded neurons. (**c**) Distribution of Pearson correlation coefficients computed over all of the neuron pairs for hS3 (1-ms time bins). *N* corresponds to the number of neuron pairs. Each of these correlation coefficients corresponds to the extremum value of the cross-correlation function computed for time lags in the interval ± 500 ms. Inset, normalized absolute (extremum) correlation coefficients for all of the neuronal pairs in the ensemble from hS3 computed for spike counts in 50-ms time bins; about 90% of the pairs had a correlation value smaller than 0.06 (vertical line). (**d**) Distribution of correlation coefficients computed over all of the neuron pairs for mLA (1-ms time bins).

For example, it was possible to correctly predict 80% of spikes in neuron 34a (participant hS3) with a false-positive rate of less than 5% (**Fig. 2a**). Similarly, an 80% true-positive rate was obtained for neuron 16a (monkey mLA) with a less than 10% false-positive rate (**Fig. 2b**). This predictability was substantial despite the weak pair-wise correlations that we observed among all neuronal pairs in the recorded datasets for hS3 and mLA (**Fig. 2c**,**d**). We confirmed this observation in our datasets for all subjects, areas and tasks.

We obtained a more comprehensive assessment of predictive performance by computing the area under the ROC curve. The area under the curve (AUC) is a global summary statistic; that is, it depends on both the true- and false-positive rates and on all of the possible thresholds on the spiking probability. The AUC gives the probability that, when two samples are randomly drawn from the data (one containing a spike, the other not), the conditional intensity model will assign a higher probability (i.e. a higher instantaneous spiking rate) to the sample with a spike[20]. It therefore provides an assessment of the discriminatory power for predictive variables under a given model. It approaches 0.5 for a chance level predictor, that is, a predictor having false- and true-positive rates along the diagonal, and it equals 1 for a perfect predictor. For example, the AUC for neuron 34a was 0.95 (**Fig. 2a**). We further corrected the AUC by subtracting chance-level predictions estimated from the actual data and scaled it to obtain a quantity that ranged from 0 (no predictive power) to 1 (perfect prediction). We refer to this corrected and normalized AUC measure as the predictive power (see Online Methods). More than 50% of examined neurons in monkey mLA showed a predictive power higher than 0.5 (median = 0.55, range = 0.1–0.89; **Fig. 3**). Prediction in subjects mCL and hS3 showed similar performances, with medians of 0.45 and 0.49 and ranges of 0.18–0.91 and 0.23–0.97, respectively. Predictive power was lower for participant hS1 (median = 0.30, range = 014–0.69). We conjectured that this was probably a result of the smaller number of neurons ($n = 21$, $n = 22$) that were recorded and used for prediction in this participant. To test this conjecture, we computed the predictive power on the basis of several smaller ensemble subsets ($n = 22$ neurons) randomly chosen out of the 110 neurons from hS3's dataset (session 2).

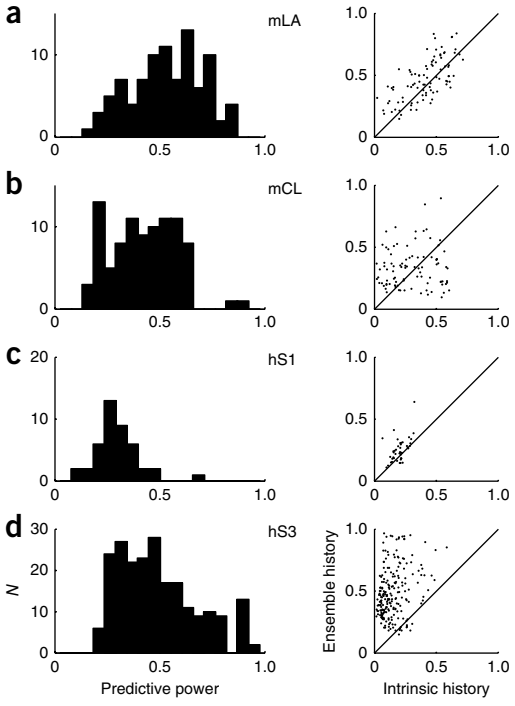

Figure 3 Predictive power of intra-areal (M1) ensemble histories. (**a**–**d**) Prediction was substantial, as shown by the distributions of predictive power corresponding to target neurons from subjects mLA (**a**), mCL (**b**), hS1 (**c**) and hS3 (**d**). Each distribution includes target neurons recorded in two different sessions (mLA: $n = 45$, $n = 45$; mCL: $n = 47$, $n = 44$; hS1: $n = 22$, $n = 21$; hS3: $n = 108$, $n = 110$). The left column shows the distribution of predictive power based on the full history model and the right column compares the predictive power of the two (intrinsic and ensemble) history components separately. The predictive power measure is based on the AUC scaled and corrected for chance level prediction. It ranges from 0 (no predictive power) to 1 (perfect prediction). For many neurons, the predictive power of separate components (intrinsic and ensemble) could add to a value larger than 1 or result in a larger predictive power than that obtained by the full history model. This indicates that there was some redundancy in the information conveyed by these two components. The numbers of predicted samples (1-ms time bins) were 864,657 and 1,230,857 for mLA, 1,220,921 and 1,361,811 for mCL, 240,000 in both sessions for hS1 and 240,000 in both sessions for hS3.

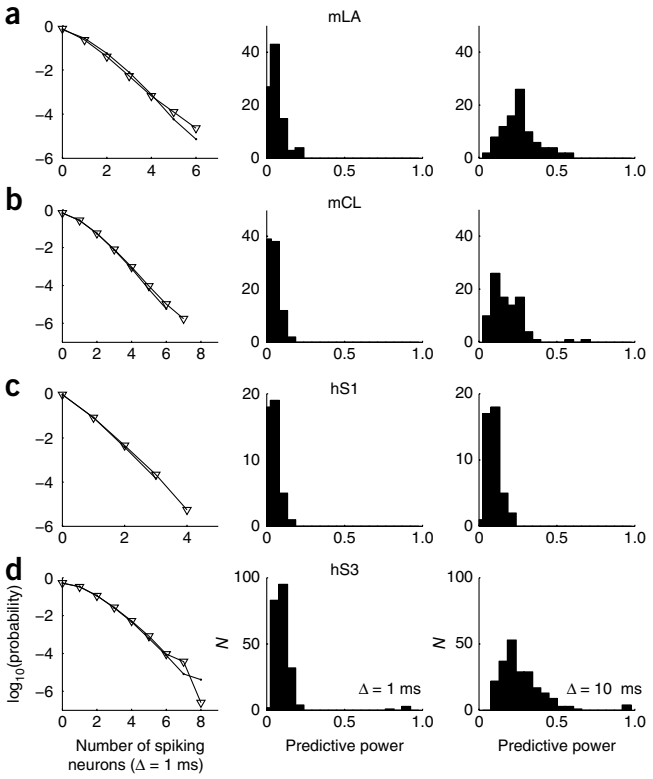

Figure 4 Predictive power of instantaneous collective states (Ising models). (**a**–**d**) The distributions of instantaneous collective states were approximated via maximum entropy distributions constrained on empirical mean spiking rates and zero time-lag pair-wise correlations. The left column shows the empirical distribution of the observed number of multi-neuron spike coincidences in the ensemble (Δ = 1 ms) and the distribution generated from the maximum entropy model via Gibbs sampling (∇). The middle column shows the distribution of predictive power values. Predictive power was computed for each target neuron separately, with the instantaneous collective state defined at a temporal resolution of 1 ms. For each given neuron, predictions were determined by a conditional spiking probability derived from the maximum entropy joint distribution model and knowledge of all of the (n – 1) neurons' simultaneous states. The right column shows the distribution of predictive power when the instantaneous state was defined at a coarser temporal resolution of 10 ms. In that case, time bins containing more than one spike were set to 1. For each monkey and human participant, data from two sessions were used in these analyses. The rows correspond to mLA (n = 45, n = 45, **a**), mCL (n = 47, n = 44, **b**), hS1 (n = 22, n = 21, **c**) and hS3 (n = 108, n = 110, **d**).

Data from hS3 provided a good reference, as neurons were recorded from M1 under the same task condition as for hS1. The distribution of predictive power values, based on this collection of smaller ensemble subsets, was similar to the one obtained for hS1, supporting our hypothesis (**Supplementary Fig. 2** and **Supplementary Note**).

These predictive power levels were obtained using the full history models, that is, the instantaneous spiking rate of each neuron conditioned on both intrinsic and ensemble histories. We further compared the predictive power of each of these two components separately (**Fig. 3**). For the majority of neurons in all of the subjects, the predictive power obtained exclusively from the ensemble histories was larger than the predictive power from intrinsic histories. In addition, it is possible that some of the intrinsic history effects, beyond refractory and recovery period effects, could also reflect network dynamics. Thus, these results indicate that the collective ensemble activity, rather than refractory and recovery dynamics, was the main source of the observed predictive power of spiking histories.

The predictive power of these M1 spiking history models in monkeys mLA and mCL was also higher than the predictive power of relevant kinematic covariates, such as hand position and velocity (**Supplementary Fig. 3** and **Supplementary Note**). We used pathlet models[21] to estimate spiking probabilities conditioned on kinematics; that is, the instantaneous spiking rate was expressed as function of hand position and velocity trajectory in the time interval (−200 ms, 300 ms) with respect to spiking time. Furthermore, the predictive power values of spiking history and pathlet models were not strictly additive (**Supplementary Fig. 4**). This result indicates that there was some level of redundancy between the information conveyed by these two models about single-neuron spiking: some of the spiking activity that was predicted by the spiking history models could also be predicted by the examined kinematics.

Predictive power of M1 instantaneous collective states

Beyond examining the predictive power of ensemble spiking histories, we were also interested in the predictive power of instantaneous collective states in the recorded ensemble, that is, simultaneous states at either 1- or 10-ms temporal resolution. Instantaneously correlated states could result from common inputs and/or from synchronization patterns arising from the neuronal network's own dynamics. Strong instantaneous collective states[7] could still be consistent with the weak pair-wise instantaneous interactions that we observed in our datasets (**Fig. 2c**,**d**). We estimated the joint probability distribution of these instantaneous collective states by fitting maximum entropy[22,23] point process models constrained by empirical mean spiking rates and zero time-lag pair-wise correlations. That is, these maximum entropy distribution models were consistent with the observed mean spiking rates and zero time-lag pair-wise correlations, but made no additional assumptions (see Online Methods). This pair-wise maximum entropy probability model, also known as the Ising model in statistical mechanics[24,25], has been shown to capture most of the instantaneous interdependency structure in *ex vivo*/*in vitro* neuronal ensemble preparations[7,8,10].

Besides satisfying the pair-wise correlation structure in the examined datasets, the estimated models also accounted well for the distribution of multiple-neuron spike coincidences, as seen in the comparison of the empirical distribution and the distribution generated from the model via Gibbs sampling[26] (**Fig. 4**). Nevertheless, the predictive power of this pair-wise correlation structure was considerably poorer than the predictive power of intrinsic and ensemble spiking histories, as captured by the models (**Fig. 4**). Across the four subjects, the predictive power values of pair-wise maximum entropy models corresponded to 0.10 ± 0.17, whereas the spiking history models yielded values of 0.50 ± 0.40 (mean ± s.d.). The predictive power of spiking history models were still higher when we considered instantaneous states at a coarser 10-ms temporal resolution (**Fig. 4**), which resulted in predictive power values of 0.26 ± 0.27 (mean ± s.d.).

Intra and inter-areal spike prediction: M1↔PMv and M1↔5d

We performed the same predictive power analysis for areas PMv and 5d and compared predictions based on ensemble histories recorded from the same or a different cortical area in the connected pairs M1↔PMv and M1↔5d (**Fig. 5**). Predictive power, determined from both local and inter-areal spiking histories, corresponding to target neurons in PMv (mCO; **Fig. 5a**) and 5d (mAB; **Fig. 5c**) had medians of 0.27 and 0.32 and ranges of 0.11–0.86 and 0.13–0.62, respectively. Predictive power corresponding to target neurons in M1, based on both local and inter-areal spiking histories, had medians of 0.50 and 0.40

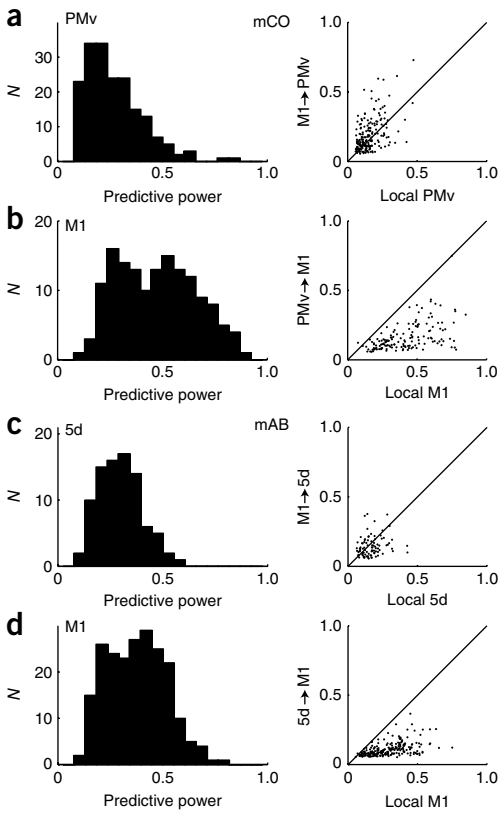

Figure 5 Predictive power of intra- and inter-areal neuronal ensemble histories. Predictive power of inter-areal ensemble history was also substantial. (**a**) Left, distribution of predictive power values for target neurons in area PMv (subject mCO), which were recorded during free reach-grasp movements. Predictive power was computed from full history models that also included spiking histories in M1. Right, comparison of the power of intra (PMv, $n = 77$, $n = 109$) and inter-areal (M1, $n = 148$, $n = 109$) ensemble histories to predict spiking in PMv. The predictive power M1→PMv tended to be higher than the local PMv→PMv in this case, where the number of neurons recorded in M1 was larger than in PMv. In contrast, additional analyses using balanced-size ensembles indicated that intra-areal predictive power was actually slightly higher (**Supplementary Fig. 5**). (**b**) Left, distribution of predictive power for target neurons in M1. Predictive power was computed from full history models that also included spiking histories in PMv. Right, comparison of the power of the intra (M1) and inter-areal (PMv) ensemble histories to predict spiking in M1. (**c**,**d**) Data are presented as in **a** and **b**, but were computed for parietal 5d ($n = 41$, $n = 47$) and M1 ($n = 104$, $n = 110$), recorded from monkey mAB during a planar pursuit tracking task. The numbers of predicted samples were 212,028 and 99,008 for mCO (sessions 1 and 2, respectively) and 416,162 and 472,484 for mAB (sessions 1 and 2, respectively).

and ranges of 0.15–0.93 and 0.11–0.83 for monkeys mCO and mAB, respectively (**Fig. 5b**,**d**). The inter-areal predictive power could also be substantial in both directions, M1 ↔ PMv and M1 ↔ 5d (**Fig. 5**). Medians were 0.15, 0.13, 0.13 and 0.09 and ranges were 0.05–0.73, 0.06–0.75, 0.06–0.37 and 0.05–0.37 for M1→PMv, PMv→M1, M1→5d and 5d→M1, respectively.

Furthermore, the predictive power from one area to target neurons in the other, especially from M1→PMv, could be larger than the predictive power of local ensembles (**Fig. 5a**,**c**). We speculated that the higher predictive power of M1→PMv in some of these cases could simply reflect the larger number of recorded neurons in M1, rather than the relative strength of feedforward and feedback connections between these two areas. A supplementary analysis using ensembles of equal size suggested that this was the case (**Supplementary Fig. 5** and **Supplementary Note**). This analysis also indicated that intra-areal predictive power was slightly higher than inter-areal predictive power when ensembles of equal size were considered.

As shown before for subjects mLA and mCL, the predictive power of spiking history models for areas M1, 5d and PMv also tended to be higher than the predictive power of pathlet models (**Supplementary Fig. 3**). The medians and ranges of predictive power values for pathlet models corresponded to 0.25 and 0.21, and 0.01–0.77 and 0.01–0.75 for mCO and mAB, respectively.

Predictability, spiking irregularity and information rates

The predictive power of spiking history models varied broadly across the 1,187 neurons in the 12 datasets. We examined whether this variation could be easily explained by simple features of the predicted spiking activity. The predictive power of history models did not appear to depend, at least in a simple manner, on the mean spiking rate or on the level of irregularity of the predicted spiking activity (**Fig. 6a–c**). Assessments of spike train irregularity were based on the coefficient of variation[3] of the interspike time intervals. We also assessed how much information was involved in the prediction of single neuron spiking.

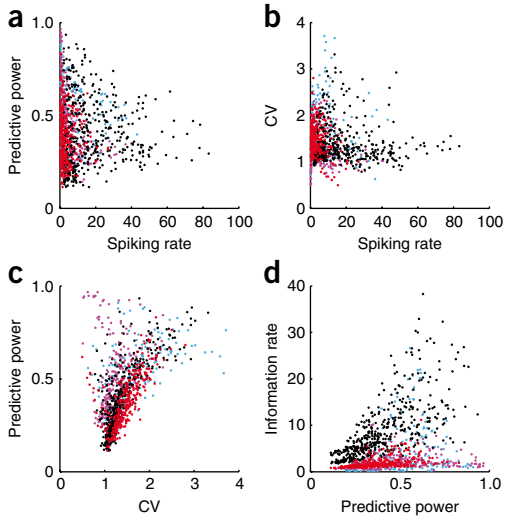

Figure 6 Predictive power, mean spiking rates, spike train irregularity and information rates. Each dot corresponds to one of the 1,187 target neurons recorded from two human participants and four monkeys, three different cortical areas and four different tasks. Each color relates to one of the different tasks. (**a**) The predictive power of full history models versus the mean spiking rate (in spikes per s) of the target neurons is shown. (**b**) Coefficient of variation (CV) of the inter-spike time intervals versus spiking rates. (**c**) Predictive power of full history models versus coefficients of variation of the predicted spiking activity. Lower coefficients indicate more regular spike trains. Coefficients around 1 and below tended to correspond to a broad range of predictive power, whereas higher coefficients tended to cluster around intermediate predictive power values. In summary, the predictive power of history models did not seem to depend, in a simple manner, on mean spiking rates or on the level of irregularity of the spiking activity. (**d**) Predictive power versus the information rate (in bits per s) involved in the prediction. Approximately equal predictive power could relate to a broad range of information rates. Blue: point-to-point reaching, monkeys mLA and mCL, area M1; purple: neural cursor control, participants hS1 and hS3, area M1; black: free reach and grasp task, monkey mCO, areas M1 and PMv; red: pursuit tracking task, monkey mAB, areas M1 and 5d.

In principle, the same level of prediction accuracy could be achieved for two different target neurons while involving different amounts of information. We examined this possibility by asking how much information about a neuron spiking or not at any given time was gained by modeling the intrinsic and ensemble history effects compared with knowing only the mean spiking rate of the process. This information gain can be estimated as the normalized difference between the log-likelihood under the estimated history model and the log-likelihood under a homogenous Poisson process with given mean spiking rate (see Online Methods). Our analysis confirmed that similar predictive power values (in the prediction of different target neurons) could involve a wide range of information rates (**Fig. 6d**).

DISCUSSION

Our findings reveal that single-neuron spiking in the cortex of humans and monkeys performing sensorimotor tasks can be substantially predicted by the spiking history of small, randomly sampled neuronal ensembles. The fact that this predictability was based on small neuronal ensembles, randomly sampled out of millions of neurons in the cortex, suggests that strongly coordinated activity underlies the generation of single-neuron spikes. This finding is notable if one considers the properties of cortical neuronal networks. Cortical neurons are embedded in large, sparsely connected recurrent networks in which the high number of synaptic inputs and high-conductance states typically induce weak coupling between randomly sampled neuron pairs. Not surprisingly, revealing and understanding these large scale and dynamic interactions has been challenging. On the basis of only the weak pair-wise correlations observed amongst cortical neurons in our datasets, one would have underestimated the strength of the statistical interdependencies induced by the collective dynamics. Furthermore, we believe our estimates of predictive power for these small neuronal ensembles should be taken as lower-bound values. There are at least a few factors that could have diminished predictive performance. For example, even though we were careful to include only trials or time segments in which the hand was moving or when the human participants were controlling a computer cursor, it is possible that the network's functional connectivity was non-stationary within and across trials. Also, we avoided more complex, and thus potentially more predictive, point process history models for computational tractability.

The fact that spiking was better predicted by the ensemble's history than by the ensemble's simultaneous collective state, estimated via pair-wise maximum entropy models, emphasizes the temporal dynamics leading to spiking. This finding, however, should not be taken as a limitation of pair-wise maximum entropy models. It is possible that multiple time-lag pair-wise correlation maximum entropy models[11,27] might capture most of the history effects detected in our data and therefore provide a simpler, minimal model. Our goal here was not, however, to provide such a minimal model for the temporal and instantaneous collective dynamics, but to determine the existence and strength of such dynamics in the cortex of humans and behaving monkeys. This is, nevertheless, an important issue that should be addressed.

In previous studies, we and others have shown that the spiking activity of small, simultaneously and randomly sampled neuronal ensembles in motor cortex can be used to predict (decode) subsequent complex behavioral variables such as arm kinematics[28–32]. Here, we found that the activity of these same neuronal ensembles can also be used to predict subsequent single-neuron spiking with substantial accuracy, implying the presence of strong collective dynamics in sensorimotor cortex. One may then ask how these collective dynamics relate to behavior. Although a comprehensive analysis of this problem is beyond the scope of this study, our results indicate that these collective dynamics do not simply reflect background coherent states that are completely unrelated to behavior and, conversely, do not simply reflect 'trivial common inputs', such as those usually considered in studies involving stimulus-driven activities of early sensory neurons. Regarding the background coherent states, our data indicate that the predictive power of models based only on ensemble history and of models based only on kinematics are not strictly additive (**Supplementary Fig. 4**). In other words, there is some level of redundancy between the information (about single-neuron spiking) conveyed by ensemble spiking histories and information conveyed by kinematics variables. Therefore, the detected collective dynamics cannot simply reflect background coherent states that are entirely unrelated to behavior. Nonetheless, it is still possible, if not likely, that part of the detected collective dynamics may reflect ongoing internal processes that are not related to behavioral variables. Regarding common inputs, we note that kinematics and other features of voluntary movements are controlled in large part (either directly, or indirectly via spinal and muscle activations) by the coordinated activity of sensorimotor cortical neurons. It is, therefore, not surprising that these behavioral variables can also achieve substantial explanatory power for spiking activity. Even so, the fact that kinematics (pathlet) models were less predictive than ensemble history models suggests that the latter carried extra information about single-neuron spiking (**Supplementary Fig. 3**). Given these considerations, the detected collective dynamics are unlikely to be explained as the simple reflection of trivial common inputs. In sum, we believe that these strong collective dynamics reflect the intra and inter-areal coordinated activity of neuronal ensembles distributed in the many different cortical and subcortical areas that participate in voluntary control of movement.

We hypothesize that the detected collective dynamics and ensemble influences on spiking activity reflect information transfer and computation in cortical networks. Collective dynamics and functional connectivity, as captured by connectivity matrices derived from ensemble history models, as well as predictive power levels, should vary as information transfer and computation change during behaviors that engage cortical areas. On the basis of current computational theories of motor control[14,15], one could predict, for example, that M1↔5d spiking predictability will manifest primarily during the initial phase of reaching movements, whereas M1↔PMv spiking predictability will peak during the hand-shaping phase of object grasping. A related and more general inquiry will be to examine the relationship of collective dynamics at this 'microscopic' spatial scale to neural activities reflected in meso- and macroscopic scale signals, such as local field potentials and electrocorticograms.

Our results also have implications for neural decoding theory and intracortical neural interfaces for motor prostheses. Collective dynamics add redundancy and, therefore, error-correcting properties to neural codes[7]. In addition, these dynamics might also account for variability of neural responses[5], which is otherwise usually attributed to noise. Therefore, it seems that ensemble history effects should be taken into consideration when decoding kinematics (or other variables) from the spiking activity of neuronal ensembles. One would predict that decoding algorithms that take into account ensemble spiking histories will outperform algorithms that treat spiking activity of different neurons as, conditioned on decoded variables, independent processes.

Our findings suggest the presence of strong collective dynamics that are fundamental to the experimental endeavor of determining coordinated spiking in cortical networks. Networks responsible for specialized cortical function are likely to be contained in the spiking

patterns of millions of neurons distributed across multiple cortical areas. Current and developing recording technologies measure the spiking activity of hundreds to thousands of neurons, a very small fraction of these networks. Without strong collective dynamics (that is, if neurons in small randomly sampled ensembles behaved seemingly independently), there would be little hope of determining how the coordinated propagation of action potentials in large-scale recurrent networks leads to computation and information processing. We believe, therefore, that the existence of these collective dynamics offers a basis for understanding cognition and adaptive behavior at the level of coordinated spiking in cortical networks.

METHODS

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

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank M.R. Fellows, C. Vargas-Irwin and B. Philip for collecting the nonhuman primate data. We thank our clinical trial participants for their dedication to this research, G. Friehs for his role as surgical investigator for the pilot clinical trial, J. Mukand for his role as clinical investigator for the pilot clinical trial, and A. Caplan, M. Serruya, M. Saleh and other employees of Cyberkinetics Neurotechnology Systems for data collection, manufacturing and clinical trial management. This study was based on work supported in part by the Office of Research and Development, Rehabilitation R&D Service, Department of Veterans Affairs (L.R.H. and J.P.D.). This work was supported by the National Institute of Neurological Disorders and Stroke (5K01NS057389-02 to W.T. and NS-25074 (Javits Award) to J.P.D.), the National Institute of Child Health and Human Development/National Center for Medical Rehabilitation Research (N01-HD-53403, subcontract to L.R.H.), the Massachusetts General Hospital Deane Institute (L.R.H.), the Doris Duke Charitable Foundation (L.R.H) and the National Institute on Deafness and Other Communication Disorders (R01-DC-009899 to L.R.H.).

AUTHOR CONTRIBUTIONS

W.T. conceived the study's central ideas and conducted the data analyses. W.T. wrote the paper with contributions from L.R.H. and J.P.D. L.R.H. and J.P.D. contributed to the clinical research design. L.R.H. was the principal investigator for the pilot clinical trial. J.P.D. supervised the nonhuman primate experiments.

COMPETING INTERESTS STATEMENT

The authors declare competing financial interests: details accompany the full-text HTML version of the paper at www.nature.com/natureneuroscience/.

Published online at http://www.nature.com/natureneuroscience/.
Reprints and permissions information is available online at http://www.nature.com/reprintsandpermissions/.

1. Braitenberg, V. & Schüz, A. *Cortex: Statistics and Geometry of Neuronal Connectivity* (Springer-Verlag, New York, 1998).
2. Elston, G.N. & Rockland, K.S. The pyramidal cell of the sensorimotor cortex of the macaque monkey: phenotypic variation. *Cereb. Cortex* **12**, 1071–1078 (2002).
3. Dayan, P. & Abbott, L.F. *Theoretical Neuroscience* (MIT Press, Cambridge, Massachusetts, 2001).
4. Destexhe, A., Rudolph, M. & Paré, D. The high-conductance state of neocortical neurons *in vivo*. *Nat. Rev. Neurosci.* **4**, 739–751 (2003).
5. Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* **273**, 1868–1871 (1996).
6. Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G. & Buzsaki, G. Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).
7. Schneidman, E., Berry, M., Segev, R. & Bialek, W. Weak pair-wise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).
8. Shlens, J. *et al.* The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.* **26**, 8254–8266 (2006).
9. Pillow, J.W. *et al.* Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
10. Yu, S., Huang, D., Singer, W. & Nikolic, D. A small world of neuronal synchrony. *Cereb. Cortex* **18**, 2891–2901 (2008).
11. Tang, A. *et al.* A maximum entropy model applied to spatial and temporal correlations from cortical networks *in vitro*. *J. Neurosci.* **28**, 505–518 (2008).
12. Jones, E.G., Coulter, J.D. & Hendry, S.H. Intracortical connectivity of architectonic fields in the somatic sensory, motor and parietal cortex of monkeys. *J. Comp. Neurol.* **181**, 291–347 (1978).
13. Shimazu, H., Maier, M.A., Cerri, G., Kirkwood, P.A. & Lemon, R.N. Macaque ventral premotor cortex exerts powerful facilitation of motor cortex outputs to upper limb motoneurons. *J. Neurosci.* **24**, 1200–1211 (2004).
14. Kalaska, J.F., Scott, S.H., Cisek, P. & Sergio, L.E. Cortical control of reaching movements. *Curr. Opin. Neurobiol.* **7**, 849–859 (1997).
15. Shadmehr, R. & Krakauer, J.W. A computational neuroanatomy for motor control. *Exp. Brain Res.* **185**, 359–381 (2008).
16. 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 (2005).
17. Paninski, L., Pillow, J. & Lewi, J. Statistical models for neural encoding, decoding, and optimal stimulus design. *Prog. Brain Res.* **165**, 493–507 (2007).
18. Truccolo, W. Stochastic point process models for multivariate neuronal spike train data: collective dynamics and neural decoding. in *Analysis of Parallel Spike Train Data* (eds. Grün, S. & Rotter, S.) (Springer, New York, in the press).
19. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **27**, 861–874 (2006).
20. Bamber, D. The area above the ordinal dominance graph and the area below the receiver operating characteristic graph. *J. Math. Psychol.* **12**, 387–415 (1975).
21. Hatsopoulos, N.G., Xu, Q. & Amit, Y. Encoding of movement fragments in the motor cortex. *J. Neurosci.* **27**, 5105–5114 (2007).
22. Jaynes, E.T. On the rationale of maximum entropy methods. *Proc. IEEE* **70**, 939–952 (1982).
23. Cover, T.M. & Thomas, J.A. *Elements of Information Theory* (Wiley & Sons, New York, 1991).
24. Landau, L.D. & Lifshitz, E.M. *Statistical Physics* 3rd edn (Butterworth & Heinemann, Oxford, UK, 1980).
25. Amit, D.J. *Modeling Brain Function* (Cambridge University Press, Cambridge, UK, 1989).
26. Geman, S. & Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721–741 (1984).
27. Marre, O., El Boustani, S., Fregnac, Y. & Destexhe, A. Prediction of spatiotemporal patterns of neural activity from pair-wise correlations. *Phys. Rev. Lett.* **102**, 138101 (2009).
28. Paninski, L., Fellows, M.R., Hatsopoulos, N.G. & Donoghue, J.P. Spatiotemporal tuning of motor neurons for hand position and velocity. *J. Neurophysiol.* **91**, 515–532 (2004).
29. Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* **408**, 361–365 (2000).
30. 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).
31. Hochberg, L.R. *et al.* Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* **442**, 164–171 (2006).
32. Truccolo, W., Friehs, G.M., Donoghue, J.P. & Hochberg, L.R. Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *J. Neurosci.* **28**, 1163–1178 (2008).

ONLINE METHODS

Human participants, electrophysiology and behavioral tasks. An investigational device exemption for these studies was obtained from the US Food and Drug Administration and all studies were performed with approval from the Spaulding Rehabilitation Hospital Institutional Review Board, the New England Institutional Review Board, the Rhode Island Hospital Institutional Review Board and the Partners Human Research Committee. The recording device, preamplifiers, data acquisition systems and computer are part of the BrainGate Neural Interface System (Cyberkinetics Neurotechnology Systems; CAUTION: investigational device. Limited by federal law to investigational use). The sensor is a 10×10 array of silicon microelectrodes that protrude 1 mm (hS1) or 1.5 mm (hS3) from a 4.2×4.2 mm platform (**Supplementary Fig. 1**). For signal acquisition, 96 electrodes are available, with minimum inter-electrode distance of 400 μm. Participant 1 (hS1) was a 24-year-old male with tetraplegia (C4 ASIA A). Participant 3 (hS3) in the pilot clinical trial is a 56-year-old female who sustained a pontine stroke 9 years before trial enrollment, resulting in loss of speech and locked-in syndrome, which later resolved to incomplete tetraplegia. After obtaining informed consent and carrying out the medical and surgical screening procedures, the array was implanted in the dominant M1 hand/arm area, identified anatomically as the 'knob' region[31,33] of the precentral gyrus in pre-operative magnetic resonance imaging.

Each participant used M1 spiking activity to control a cursor displayed on a computer screen[31,32]. The participant was instructed to imagine moving a circular cursor displayed on the screen to one of four peripheral targets, positioned at 0, 90, 180 and 270°. In each session, 20 trials were collected for each of the four pseudo-randomly presented radial targets. The 3-s period after target appearance was included in the analysis. Two datasets per participant were used, corresponding to two research sessions conducted on two different days. Data sets were collected on 2004.09.16 and 2004.09.20 (hS1) and 2006.01.23 and 2006.01.24 (hS3). We used spike-sorting utilities in Offline Sorter (Plexon) to identify and sort neuronal units in all of the human and nonhuman primate recordings. We did not distinguish whether a single unit sorted from the same electrode on different days corresponded to the same neuron or not.

Nonhuman primate subjects, electrophysiology and behavioral tasks. Two datasets were recorded in two experimental sessions from each of four rhesus monkeys (*Macaca mulatta*). All recordings were obtained via single or dual cortically implanted 10×10 microelectrode arrays (electrode length, 1.0 mm), similar to the array described above. M1 neurons from monkeys mLA and mCL were recorded while they performed point-to-point planar movements. The monkeys used a manipulandum to move a position feedback cursor that was presented on the monitor. Targets were randomly placed (that is, uniform in two-dimensional space), one at a time, on the workspace. After the successful acquisition of a random number (3–9) of targets, the monkeys received a juice reward. Only segments of data recorded during the reaching phases of the tasks, from two experimental sessions, were included in the analyses (datasets: mLA, 2004.03.25 and 2004.03.26; mCL, 2004.03.25 and 2004.03.29). M1 and PMv neurons were simultaneously recorded from monkey mCO while this monkey performed reaching and grasping movements toward objects moving in the workspace (C.E. Vargas-Irwin, P. Yadollahpour, G. Shakhnarovich, M.J. Black and J.P. Donoghue, *Soc. Neurosci. Abstr.* 673.18, 2008). A motion-capture system was used to record arm-hand configurations and related behavioral epochs. Only data segments corresponding to 1-s segments during the reach-grasp phase before the final object grasp, from two sessions, were included in the analyses (datasets: 2008.03.19 and 2007.12.12). M1 and 5d neurons were simultaneously recorded via dual arrays from monkey mAB. The monkey performed visually guided pursuit tracking of a circular cursor projected on a horizontal screen while wearing an external device for kinematic measurements (Kinarm, BKIN Technologies). This cursor followed randomly generated trajectories of varying speeds over the planar, horizontal workspace (B.A. Philip and J.P. Donoghue, *Soc. Neurosci. Abstr.* 672.22, 2008). Only data segments recorded during tracking, from two sessions, were included in the analyses (datasets 2008.05.08 and 2008.05.09). All procedures were in accordance with Brown University Institutional Animal Care and Use Committee approved protocols and the Guide for the Care and Use of Laboratory Animals.

Point process history models. The distribution of the point process sample paths of a given i^{th} neuron is completely specified[16,18,34] by the conditional intensity function (instantaneous conditional spiking rate)

$$\lambda_i(t \mid \mathcal{H}(t), z(t)) = \lim_{\Delta \to 0} \frac{\Pr(N_i(t+\Delta) - N_i(t) = 1 \mid \mathcal{H}(t), z(t))}{\Delta} \quad (3)$$

where $\Pr(\cdot|\cdot)$ is a conditional probability, $N_i(t)$ denotes the sample path (that is, a right-continuous function that jumps 1 each time a spike occurs), $\mathcal{H}(t)$ denotes the conditioning intrinsic and ensemble spiking histories up to, but not including, time t, and $z(t)$ denotes other relevant extrinsic covariates, such as stimuli and behavioral variables. We focused on intensity function models conditioned on spiking histories (see **Supplementary Fig. 3** and **Supplementary Note** for analyses involving conditional intensity function models that also included extrinsic covariates such as hand position and velocity). The sample path distribution for the discrete time point process belongs to the exponential family with canonical parameter $\log(\lambda(t|\mathcal{H}_t)\Delta)$, which we modeled as

$$\log(\lambda_i(t \mid \mathcal{H}_t)\Delta) = \mu_i + \sum_{k=1}^{10} K_{1,i,k} \cdot x_i + \sum_{j \neq i} \sum_{k=1}^{4} K_{2,i,j,k} \cdot x_j \quad (4)$$

where t indexes discrete time, Δ=1 ms, μ_i relates to a background level of spiking activity, and x_i denotes the spiking history (spike train) in the time interval (t−100 ms, t) for the i^{th} neuron, with $x_{i,t} \in \{0,1\}$ for i=1,2,…,n recorded neurons. $K_{1,i,k}$ and $K_{2,i,j,k}$ consisted of temporal basis functions of the raised cosine type[9] with coefficients to be estimated. Ten and four basis functions were used for the intrinsic and ensemble history filters, respectively. Thus, $K_{1,i}$ and $K_{2,i,j}$ in equation (2) consisted of nonparametric temporal filters for the intrinsic and ensemble spiking histories, respectively. Consistent with known neurophysiology and measured autocorrelation functions of the recorded spike trains, we enforced an absolute refractory period of 2 ms in the intrinsic history component. A history model for a particular neuron did not include the spiking history of other neurons recorded by the same electrode. That was done to avoid potential negative correlation artifacts, especially at zero and short time lags, commonly introduced by current spike thresholding and sorting algorithms. This rule was also adopted in the computation of the distribution of pair-wise correlation coefficients. Model parameters were estimated via gradient-ascent maximization of the penalized log-likelihood functions[18]. A regularization term in the form of a ridge regression penalty was added for the model parameters related to the ensemble history effects. After estimating a conditional intensity function model, the probability of a given neuron spiking at any given 1-ms time bin, conditioned on past spiking histories, was obtained as[16]

$$\Pr(x_{i,t} = 1 \mid \mathcal{H}_t) = \lambda_i(t \mid \mathcal{H}_t)\Delta + o(\Delta) \approx \lambda_i(t \mid \mathcal{H}_t)\Delta \quad (5)$$

The term $o(\Delta)$ relates to the probability of observing more than one spike in a 1-ms time interval.

Instantaneous collective states: pair-wise maximum-entropy point process models. The total interdependence in multivariate stochastic processes can be decomposed into two main components[35]: a time 'causal' component (that is, the statistical dependence of current states on past events) and an instantaneous component (for example, instantaneous dependencies among neurons). We considered instantaneity at two time resolutions: 1- and 10-ms time bins. When using the 10-ms resolution, the rare cases of time bins with more than one spike were represented as 1-spike events. We estimated statistical interdependencies in these instantaneous collective states by first fitting zero time-lag pair-wise maximum entropy distribution models[7,22]. Estimation of probability distributions for high-dimensional systems without further constraints is typically an intractable problem. On the other hand, second-order statistics in the form of pair-wise correlations are still feasible to compute and maximum entropy distributions constrained on pair-wise correlations can then be estimated. This maximum entropy distribution model constrained on mean rates and pair-wise zero time-lag correlations[7,23] is given by

$$\Pr(x_{1,t}, x_{2,t}, ..., x_{n,t}) = \frac{1}{Z_{(\alpha,\beta)}} \exp\left\{ \sum_i \alpha_i x_{i,t} + \frac{1}{2} \sum_{i \neq j} \beta_{i,j} x_{i,t} x_{j,t} \right\} \quad (6)$$

where $x_{i,t} \in \{-1,1\}$, corresponding to no spike and spike, respectively, $Z_{(\alpha,\beta)}$ is a normalization constant, $\{\alpha_i\}$ reflects constraints imposed by the empirical mean

doi:10.1038/nn.2455

spiking rates, and $\{\beta_{ij}\}$ reflects constraints imposed by the zero time-lag pair-wise correlations, with $\beta_{ij}=\beta_{ji}$. The conditional spiking probability of a given neuron under this pair-wise maximum entropy distribution model is given by

$$\Pr(x_{i,t}=1 \mid x_{\setminus i,t}) = \frac{1}{2} \frac{\exp\left\{\alpha_i + \sum_{j\neq i} \beta_{i,j} x_{j,t}\right\}}{\exp\left\{\alpha_i + \sum_{j\neq i} \beta_{i,j} x_{j,t}\right\} + \exp\left\{-\alpha_i - \sum_{j\neq i} \beta_{i,j} x_{j,t}\right\}} \quad (7)$$

where $x_{\setminus i,t}$ denotes the observed neuronal ensemble state not including the i^{th} neuron. Parameters of these pair-wise maximum entropy models were estimated via maximum pseudolikelihood[36]. The consistency of this estimator has been previously demonstrated[37–39], as has its relationship to contrastive divergence methods[39]. We used Gibbs sampling[26] to sample from the estimated maximum entropy models and compute the distributions of multiple-neuron spike coincidences in 1- and 10-ms time bins. For practical reasons, and in contrast with the ensemble history point process models, the data used to fit the maximum entropy distribution models also included neurons isolated in the same recording channel. This could have, if anything, helped to improve the performance of these maximum entropy models in comparison with the ensemble history models.

ROC analysis. ROC curves are a standard tool for the analysis of prediction performance[19]. After estimating a conditional intensity function model on training data, the probability of a given neuron spiking at any given 1-ms time bin, conditioned on either past spiking histories or instantaneous states, was computed on test data according to equations (5) and (7), respectively. A tenfold-crossvalidation scheme was used. From this probability, true- and false-positive prediction rates were computed, resulting in the ROC curves. We used the AUC to derive a predictive power measure. Informally, the relationship between AUC and predictive power can be expressed as follows. Consider the case where all of the 1-ms time bin samples are separated into two populations, one consisting of samples with a spike and the other consisting of samples with no spikes. Next, consider randomly drawing two samples, one from each population. The AUC gives the probability that our conditional intensity function model will assign a higher probability (that is, a higher instantaneous spiking rate) to the sample from the spike population than to the sample from the no-spike population[20]. The AUC therefore provides an assessment of the discriminatory or predictive power for predictive variables under a given model. It also relates to the Wilcoxon-Mann-Whitney U statistics in the case of independent samples. Given the temporal dependencies in our data, we computed the AUC directly from the ROC curve and used random permutation tests to establish the statistical significance of estimated values. Typical confidence intervals were extremely narrow overall. The ROC curve of a chance level predictor asymptotes the diagonal line, resulting in a AUC of 0.5. In our datasets, the AUC corresponding to chance prediction could be slightly larger than 0.5. We defined a predictive power, with respect to this chance level, as $2\times(AUC-AUC^*)$, where AUC^* is the AUC corresponding to a chance level predictor for a particular neuron and model, estimated via random permutation methods. The scaling by a factor of 2 was introduced so that the predictive power ranged from 0 (no predictive power) to 1 (perfect prediction).

Information rates. We computed an information rate that estimates how much reduction in the uncertainty about whether or not a neuron will spike in any given time bin, conditioned on knowing only the mean spiking rate, is achieved by also knowing spiking histories and their estimated effects under a given history model. Given large enough number of samples and under an ergodicity assumption, this information rate can be approximated as[23]

$$\frac{1}{T}\sum_{t=1}^{T} x_{i,t}\log(\lambda_i(t \mid \mathcal{H}_t)\Delta) - \lambda_i(t \mid \mathcal{H}_t)\Delta - (\bar{\lambda}_i\Delta\log(\bar{\lambda}_i\Delta) - \bar{\lambda}_i\Delta) \quad (8)$$

where $x_{i,t}\in\{0,1\}$, $\bar{\lambda}_i$ is the mean spiking rate and $\lambda_{i,t}=\lambda_i(t|\mathcal{H}_t)$. The summation term corresponds to the average log-likelihood under the history model, with averages computed over T samples (1-ms time bins) and the second term corresponds to the average log-likelihood under a homogeneous Poisson process with the specified mean rate. Computed to base 2 and normalized by Δ, the above quantity corresponds to an information rate in bits per second.

33. Yousry, T.A. *et al.* Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. *Brain* **120**, 141–157 (1997).
34. Daley, D.J. & Vere-Jones, D. *An Introduction to Theory of Point Processes* (Springer, New York, 2003).
35. Geweke, J. Measurement of linear dependence and feedback between multiple time series. *J. Am. Stat. Assoc.* **77**, 304–314 (1982).
36. Besag, J. Statistical analysis of non-lattice data. *Statistician* **24**, 179–195 (1975).
37. Geman, S. & Graffigne, C. Markov random field image models and their application to computer vision. *Proc. Int. Congr. Math.* 1496–1517 (1986).
38. Gidas, B. Consistency of maximum likelihood and pseudo-likelihood estimators for Gibbsian distributions. in *Stochastic Differential Systems, Stochastic Control Theory and Applications* (eds. Fleming, W. & Lions, P.-L.) (Springer, New York, 1988).
39. Hyvärinen, A. Consistency of pseudolikelihood estimation of fully visible Boltzmann machines. *Neural Comput.* **18**, 2283–2292 (2006).

Collective Dynamics in Human and Monkey Sensorimotor Cortex: Predicting Single Neuron Spikes

Supplementary Information

Wilson Truccolo[1,2,5], Leigh R. Hochberg[2-6] and John P. Donoghue[4,1,2]

[1]Department of Neuroscience, [2]Brown Institute for Brain Science and [3]Division of Engineering, Brown University, Providence, RI
[4]Rehabilitation Research and Development Service, Department of Veterans Affairs, Providence, RI
Departments of Neurology, [5]Massachusetts General, [6]Brigham and Women's, and Spaulding Rehabilitation Hospitals; Harvard Medical School, Boston, MA
Corresponding author: wilson_truccolo@brown.edu

Supplementary Note

A. Single and dual microelectrode array recordings

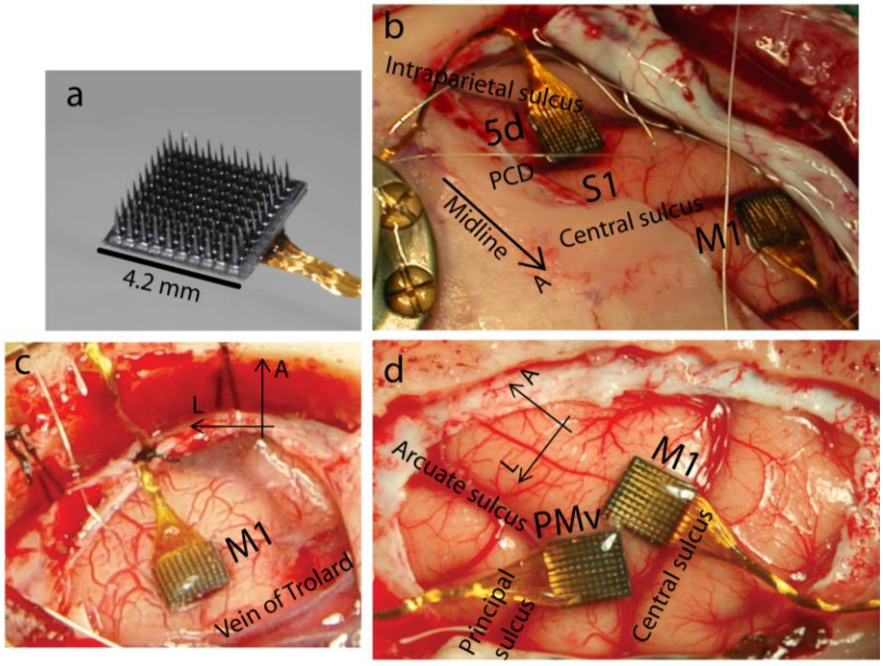

Supplementary Figure 1. Single and dual microelectrode array recordings from human and non-human primate sensorimotor cortex. Chronically implanted arrays are shown during the surgery. Recordings reported here were performed weeks to months after surgical implantation. **(a)** The 10 x 10 microelectrode array. The array's platform is 4.2 x 4.2 mm, with minimum inter-electrode distance of 400 μm. Maximum inter-electrode distance was ~ 2 cm (for electrodes in two arrays). **(b)** Dual array recordings: two implanted arrays in arm related areas of monkey primary motor (M1) and parietal (5d) cortex. PCD stands for posterior central dimple and "midline" corresponds to the sagittal suture. The arrow point to the anterior (A) direction. **(c)** Array implanted in the arm (knob) area of primary motor cortex, human subject hS3. The labeled vein of Trolard is a large superficial vein that runs atop the central sulcus. The arrows point to anterior and lateral (L) directions. **(d)** Dual array recordings from monkey M1 and ventral premotor (PMv) areas. The twelve datasets used in the analyses included 1,187 neuronal recordings: hS1 (n = 22, n = 21), hS3 (n = 108, n = 110), mLA (n = 45, n = 45), mCL (n = 47, session 1; n = 44, session 2), mCO (M1: n = 148, n = 109;

Nature Neuroscience: doi:10.1038/nn.2455

PMv: n = 77, n = 109) and mAB (M1: n = 104, n = 110; 5d: n = 41, n = 47). We did not distinguish whether a single unit sorted from the same electrode on different days corresponded to the same neuron or not.

B. Predictive power and ensemble size in M1

We conjectured (main text) that the smaller size of the M1 neuronal ensemble for participant hS1 (n=21, n=22) explained the lower predictive power obtained for this participant in comparison to hS3, mLA and mCL (Fig. 3, main text). To further investigate this possibility, we randomly sampled subsets of 22 neurons out of the 110 neuron ensemble from participant hS3 (session 2). Data from hS3 provided a good reference since these neurons were recorded from M1 under the same task condition. Twenty of these random subsets, each of size n = 22, were sampled. Although these subsets were different, some of the neurons could be present in multiple subsets. In each random subset, conditional intensity functions for each neuron were modeled as a function of the intrinsic spiking history and the spiking histories of the other neurons in the random subset. Predictive power of these history models for each target neuron was computed as in the main text. Supplementary Figure 2 shows the distribution of predictive power for all of the neurons in these 20 random subsets. This distribution was similar to the one obtained for hS1, supporting our conjecture.

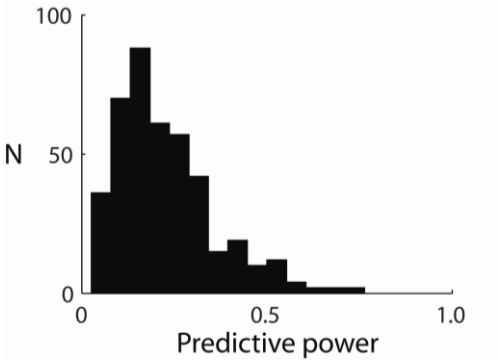

Supplementary Figure 2. Distribution of M1 (Intra-areal) ensemble predictive power: random ensemble subsets of size n = 22 (hS3, session 2).

C. Comparing the predictive power of spiking history and pathlet models (hand position and velocity trajectories) in M1, PMv and 5d neurons

M1 neurons are known to be strongly modulated with hand position and velocity during reaching movements, as performed in the task executed by the monkeys in this study (Ashe and Georgopoulos, 1994; Moran and Schwartz, 2001). In particular, Hatsopoulos et al. (2007) and Paninski et al. (2004) have shown that M1 neurons can be highly tuned to preferred trajectories in velocity space, that is, neurons are not only tuned to the instantaneous velocity at a particular single time lag, but are related to velocity trajectories spanning a short time period. This means that trajectories (histories) in velocity space can also predict single neuron spikes and could potentially account for the same predictive power observed for the ensemble history. We were interested in how the predictive power of recent spiking history compared to the predictive power of hand position and velocity trajectories. We used 'pathlet' models (Hatsopoulos et

Nature Neuroscience: doi:10.1038/nn.2455

al., 2007) to express the instantaneous spiking rate at time *t* as a function of the instantaneous hand position and velocities over the time interval [-200, 300] ms around neuron time *t*. Specifically, this 'pathlet' conditional intensity model was obtained from

$$\log\left[\lambda_i(t \mid p_{1,t+\tau}, p_{2,t+\tau}, v_{1,t+\tau_1:t+\tau_2}, v_{2,t+\tau_1:t+\tau_2})\Delta\right] = \mu_i + a_i p_{1,t+\tau} + b_i p_{2,t+\tau} + \sum_{k\in\{-4,-1,\ldots,6\}} c_{i,k} v_{1,t+k\delta} + d_{i,k} v_{2,t+k\delta} \ ,$$

where $p_{1,t+\tau}$ and $p_{2,t+\tau}$ are the positions in the horizontal and vertical coordinates at a preferred time lag $t+\tau$ (depending on the neuron), $v_{1,t+\tau_1:t+\tau_2}$ and $v_{2,t+\tau_1:t+\tau_2}$ correspond to the velocities over the time interval $[\tau_1 = -200, \tau_2 = 300]$ ms, sampled at $\delta = 50$ ms time steps; (pathlet models for *m*CO included position and velocity in 3 dimensions); a,*b,c* and *d* are tuning parameters estimated via maximum likelihood methods (Truccolo et al., 2005). The choice of the above time interval was based on the study in Hatsopoulos et al. (2007) and on exploration of different intervals in our own data; use of larger time intervals did not significantly improve prediction performance. Supplementary Figure 3 (top row) shows the distributions of predictive power values based on these pathlet models for neurons recorded from the four monkeys. The predictive power of full history models was typically higher than the predictive power of pathlet models.

We also examined the issue of redundancy between the information conveyed by spiking histories and information conveyed pathlet models (Supplementary Fig. 4). The following points summarize our analysis of this redundancy:

(a) If there was no redundancy in the information conveyed by these two models, the predictive powers of a history model and of a pathlet model for a given neuron should add to at most 1, the maximum possible value corresponding to perfect prediction. By contrast, the predictive powers of these two individual models were not strictly additive, i.e. they could add up to values greater than 1, indicating that some of the same spiking activity predicted by the spiking history models could also be predicted by the examined kinematics (Supplementary Fig. 4).

(b) This redundancy was also present for neurons whose summed predictive powers of spiking history and pathlet models added up to a value smaller than 1. The predictive power of a combined model (i.e. a model that included both spiking histories and kinematics) was significantly smaller than the sum of the predictive power of the two individual models for the majority of neurons (>80% of neurons in the 3 studied cortical areas).

(c) Despite the redundancy between the information conveyed by these two models as implied in (a,b), there was also extra information about single neuron spiking in the spiking history that could not be accounted for by the examined kinematics. This follows from the fact that predictive power was typically higher for spiking history models as seen in Supplementary Figure 3.

Nature Neuroscience: doi:10.1038/nn.2455

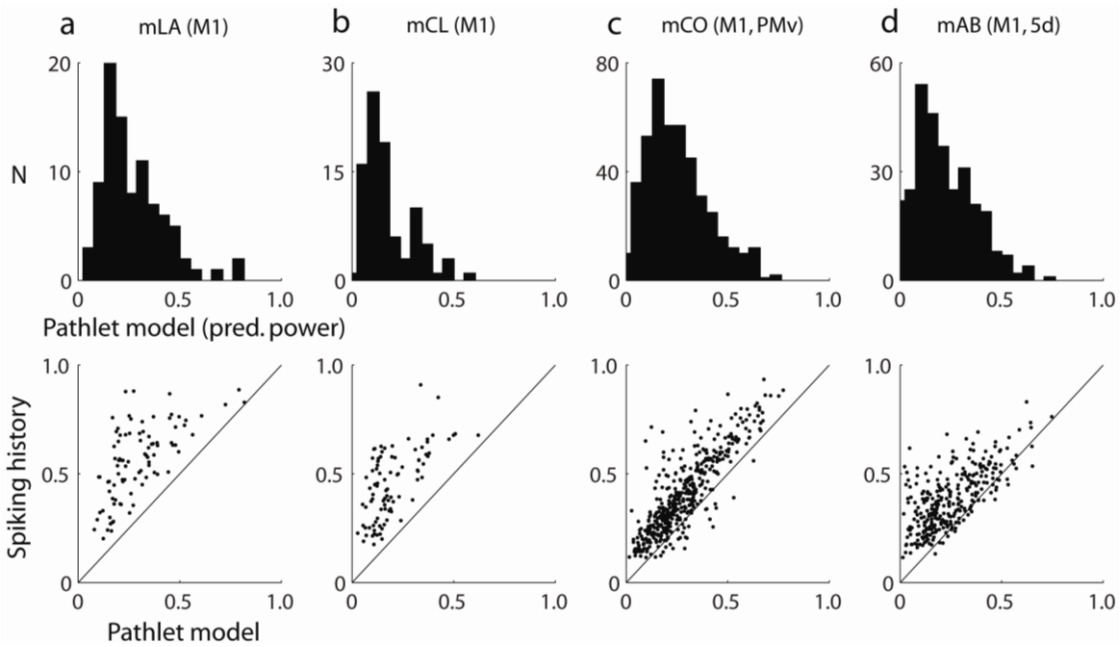

Supplementary Figure 3. Predictive power of spiking history models versus pathlet models. Top row shows the distributions of predictive power computed from the pathlet models for each of the four monkeys. The second row shows the predictive power of spiking history models versus the predictive power of pathlet models.

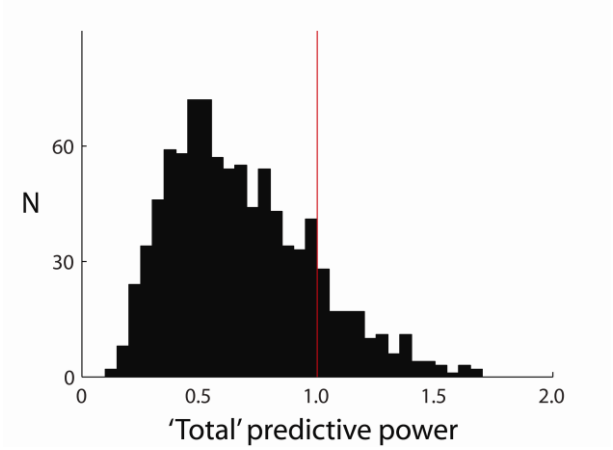

Supplementary Figure 4. Histograms of the 'total' predictive power of ensemble spiking history and pathlet models. Here, 'total' predictive power refers to the sum of the predictive power of the spiking history model and the power of the pathlet model, computed for each neuron (n = 924 neurons; datasets mLA, mCL, mCO and mAB). This 'total' predictive power should add up to at most 1 (perfect prediction) if there was no redundancy between the information conveyed by ensemble history models and the information conveyed by pathlet models.

D. Ensemble size and inter-areal predictive power

Further, we also investigated whether the larger number of neurons in M1 ensembles could explain the higher predictive power of inter-areal predictive power, in comparison to the intra-areal predictive power, especially for the M1 → PMv effect in some of the examined target neurons. To address this issue, we fitted new history models based on balanced-size ensembles, that is, the number of neurons in M1 and PMv

Nature Neuroscience: doi:10.1038/nn.2455

ensembles (and in the M1 and 5d ensembles) were set to be equal. First, we set the number of neurons *n* in the ensembles for both areas to equal the number of neurons in the smallest ensemble of the cortical pair: in our case, the number of neurons in 5d for the M1−5d pair, and PMv for the M1−PMv pair. Second, for each neuron whose spiking was to be predicted, we ranked the input neurons according to their strength based on the magnitudes of the corresponding original model coefficients: i.e., the coefficients estimated based on the full history models shown in Figure 5 (main text). We then included in the new *balanced-size* models only the *n* top-ranked neurons. Supplementary Figure 5 shows the differences between intra and inter-areal predictive power for target neurons in M1, PMv and 5d. Differences between intra and inter-areal predictive power in these balance ensembles were small on average. Nevertheless, especially for M1, the distribution of these differences was skewed toward positive values, i.e. higher values for intra-areal predictive power.

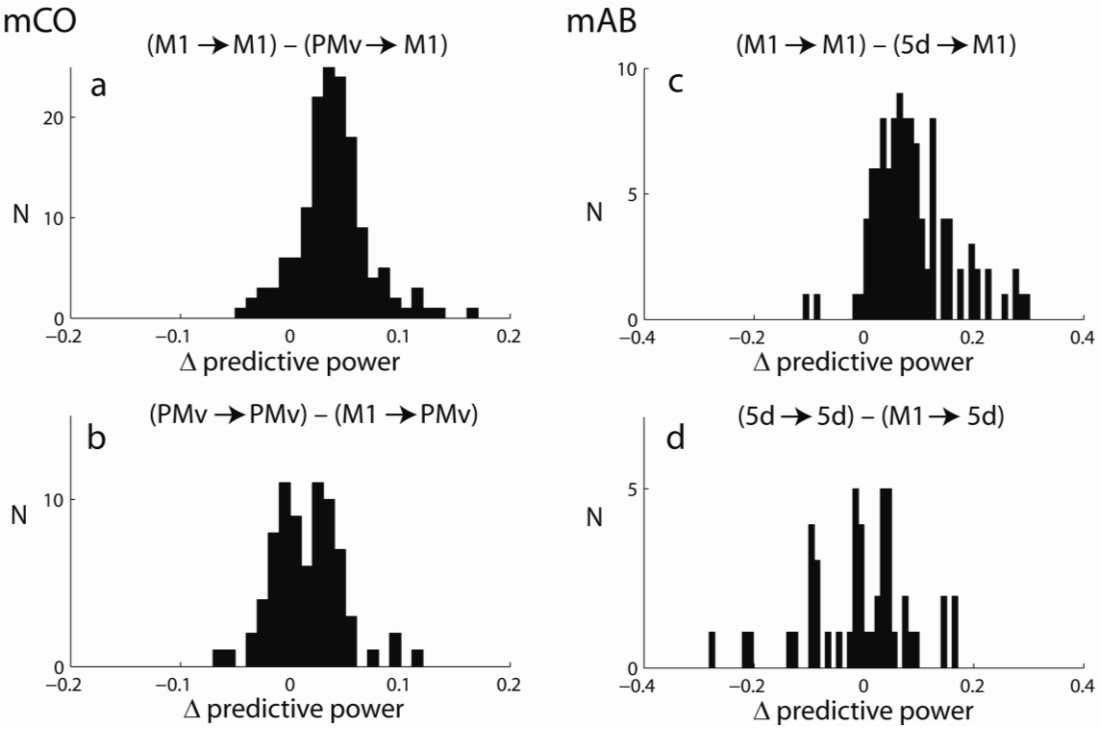

Supplementary Figure 5. Intra versus inter-areal predictive power for balanced-size ensembles. (a,b) and (**c,d**) show the comparisons for the M1−PMv and M1−5d pairs, respectively. The term "Δ rel. predictive power" denotes the difference between intra and inter-areal ensemble predictive power. For example, (PMv → PMv) − (M1 → PMv) represents the difference between the predictive power of intra-areal ensembles in PMv and inter-areal ensembles in M1 to predict single neuron spiking activity in PMv. Balanced-size ensembles: M1−PMv, n = 77 and n = 109, for sessions 1 and 2, respectively; M1−5d, n = 41 for both sessions.

Supplementary references

Ashe, J. & Georgopoulos, A.P. Movement parameters and neural activity in motor cortex and area-5. *Cereb. Cortex* **4**, 590–600 (1994).

Moran, D.W. & Schwartz, A.B. Motor cortical representation of speed and direction during reaching. *J. Neurophysiol.* **82**, 2676 –2692 (1999).

Nature Neuroscience: doi:10.1038/nn.2455

Hatsopoulos, N.G., Xu, Q. & Amit, Y. Encoding of movement fragments in the motor cortex. *J. Neurosci*. **27**(19), 5105-5114 (2007).

Paninski, L., Fellows, M. R., Hatsopoulos, N. G., & Donoghue J. P. Spatiotemporal tuning of motor neurons for hand position and velocity. *J. Neurophysiol.* **91**, 515–532 (2004).

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 (2005).

Nature Neuroscience: doi:10.1038/nn.2455

VOLUME 14 NUMBER 5 MAY 2011
www.nature.com/natureneuroscience

nature neuroscience

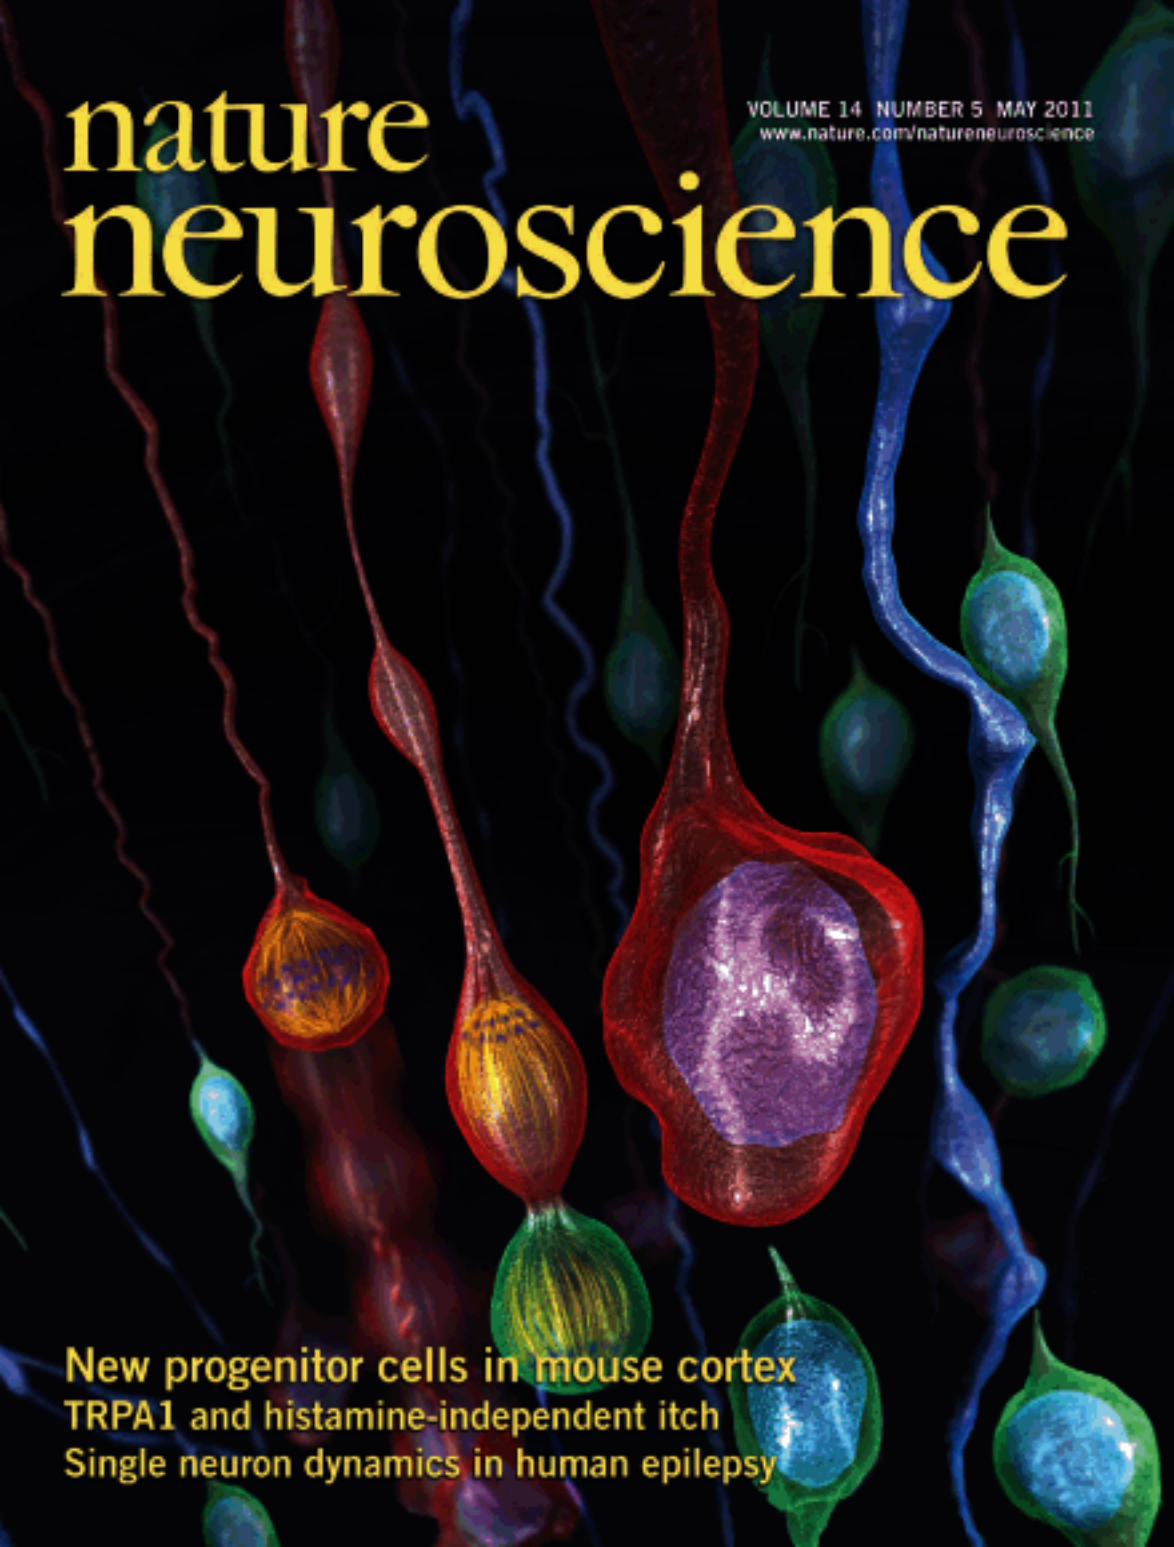

New progenitor cells in mouse cortex
TRPA1 and histamine-independent itch
Single neuron dynamics in human epilepsy

The ups and downs of seizure activity

Matthew C Walker

A study this issue by Truccolo *et al.* analyzing extended recordings of single-neuron activity in human neocortical epilepsy, demonstrates that, even in areas remote from the seizure focus, neuronal firing patterns alter minutes before seizure onset, are heterogeneous during seizures, and change homogeneously at seizure offset.

Epilepsy has always been shrouded in mystery: what is it that can rob people of their senses, yet return them almost unscathed? This striking manifestation of seizures led to ancient theories of possession by gods or demons. The modern concept of epilepsy arrived with Hughling Jackson's 1873 definition, "Epilepsy is the name for occasional, sudden, excessive, rapid, and local discharges of grey matter"[1]. In the ensuing 140 years, despite enormous advances in our understanding of the neurochemical, synaptic and ionic basis of epilepsy, our concept of what constitutes a seizure has not advanced much beyond this definition. In 1936, William Lennox said of neuronal activity during a seizure, "The harmony of the symphony orchestra has become a single note"[2], and this has remained a prevailing view. However, later animal experiments indicated that the idea of a seizure as the uniform synchronous firing of neurons needed to be modified, because individual neurons could show very different firing patterns during seizures (increasing or decreasing their firing rate)[3–5]. This heterogeneity had also been observed in single-unit neuronal recordings from the hippocampi of humans with epilepsy[6]; the number of neurons 'involved' in a seizure was suggested to predict the extent to which the seizure would spread[6].

In this issue, Truccolo and colleagues have undertaken the most comprehensive study yet of single-unit activity during human neocortical seizures, and the results provide considerable insights into the dynamics of neural activity in epilepsy[7]. Four patients who were undergoing intracranial electrode placement for treatment of their neocortical epilepsy had a 16-mm^2 grid of 100 microelectrodes positioned close to the area of seizure onset (up to 4 cm away), with the electrodes penetrating cortical layer III/IV. It could be argued that, because the electrodes are distant from the seizure onset zone, this is a study of areas of seizure propagation rather than of seizure onset, and neuronal responses during a seizure may be different in these regions[5]. This does not, however, detract from the importance of the findings, and in one case the microelectrodes were placed in the area of seizure onset. These microelectrodes were used to detect the firing of single neurons before, during and after seizures (the pre-ictal, ictal and post-ictal periods, respectively), and eight seizures were selected, on the basis of clearly separable single units, from among the many seizures recorded.

The number of active neurons and the change of neuronal firing rate at the beginning of the seizure varied between seizures; however, the heterogeneity of ictal spiking rate across the neuronal population consistently increased: the firing rate of individual neurons increased, remained unchanged or decreased[7] (**Fig. 1a**). Such heterogeneity of neuronal firing at the focus and in areas of seizure propagation had been described in earlier animal studies[3–5], but in some of those studies the heterogeneity could be attributed to the method of seizure generation and/or the presence of anesthesia. Indeed, similar animal studies without anesthesia did not reproduce the same degree of heterogeneity[8]. Notably, most of these studies used an acute focus (penicillin application) or electrical stimulation to induce seizures, a situation that is very different from the spontaneously occurring seizures in human epilepsy. Nevertheless, more recent animal evidence has indicated that such heterogeneity can be observed in unanesthetized animals having spontaneous seizures[9].

A critical question is what underlies this heterogeneity. It had been hypothesized that there is a sequential recruitment of neurons involving first a "depressed" state, followed by a "projected" state in which the neuron would respond to the input but inhibitory postsynaptic potentials (IPSPs) would still be present, and lastly a "propagated" state in which the neuron would be persistently depolarized and IPSPs could not be elicited[5]. The simultaneous local occurrence of different neuronal states in this study argues against this hypothesis. Nevertheless, such heterogeneity raises the question of whether the ictal firing pattern of an individual neuron is secondary to the intrinsic properties and connectivity of that neuron or is due to the degree of recruitment. Earlier animal studies had suggested that a neuron could change its behavior from one seizure to the next, supporting the latter possibility[3]. Truccolo and colleagues[7], however, make two important observations: first, by using the characteristics of the spikes to distinguish between those generated by interneurons and by pyramidal cells, they find that this distinction cannot explain the heterogeneity of spiking patterns; and second, in a few instances, they were able to record repeated seizures in the same person occurring at short enough intervals to permit reliable reidentification of the same units and observed that neurons tended not to change their behavior from seizure to seizure. These findings indicate that the firing pattern of a neuron during a seizure is a characteristic of that neuron. Many questions remain unanswered; for example, what is the mechanism underlying differences in neuronal behavior, and are these differences secondary to intrinsic neuronal properties, local connectivity or more distributed networks?

Two further observations from this study perhaps constitute the most remarkable findings[7]. First, the end of the seizure is marked by an abrupt, homogeneous change in neuronal firing. Because this affects both interneurons

Matthew C. Walker is at the University College London Institute of Neurology, University College London, London, UK.
e-mail: mwalker@ion.ucl.ac.uk

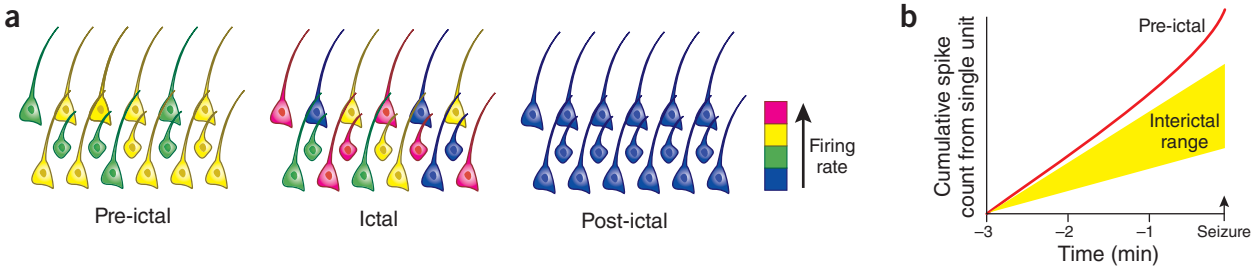

Figure 1 Changes in neuronal (unit) firing before, during and after a seizure. Pre-ictal, ictal and post-ictal activity are recorded at sites of seizure propagation. (**a**) During the ictal period, there is a variable change in the firing rate of individual neurons, leading to increased heterogeneity of the neuronal population. At seizure termination (post-ictal period), there is a sudden, homogeneous decrease in neuronal firing. (**b**) The cumulative spike counts of some neurons deviate from their 'normal' range in the minutes before a seizure occurs, indicating that a distinct pre-ictal period exists in even areas remote from the focus.

and pyramidal cells, it is unlikely to be due to local inhibitory circuits. There is evidence from animal models that, during seizures, extracellular potassium accumulation can first increase neuronal firing but then, at higher concentrations, inactivate sodium channels, resulting in a depolarization block of spike generation, termination of the seizure and post-ictal depression[10,11]. However, this mechanism is also an unlikely explanation for human neocortical seizure termination. During a depolarization block, inactivation of sodium channels would be expected to alter the spike waveform, and this was not observed during single-unit suppression at seizure termination[7]. Nevertheless, it is possible that seizure termination within the focus itself is secondary to depolarization block, and that seizure termination in areas of propagation is then due to decreased synaptic input. Indeed, a sudden increase in inhibition through extracellular increases in adenosine, GABA or other neuromodulators could also decrease synaptic input, or there could be a critical inhibitory input from more distant sources such as the basal ganglia[12]. An alternative, appealing explanation—and one that is supported by human intracranial and scalp EEG recordings—is that seizure termination is, paradoxically, secondary to increased neuronal synchronization, which then drives the network into a refractory state[13].

Second, Truccolo and colleagues[7] identified changes that occur in neuronal firing before seizure onset (**Fig. 1b**). Evidence of a pre-ictal period has previously been obtained from scalp and intracranial EEG recordings; however, methods of EEG analysis for seizure prediction have suffered because of poor specificity, sensitivity or both[14]. Trucculo and colleagues[7] observed in some units, minutes before seizure onset, a deviation in the trajectory of the cumulative spike count with time from that observed interictally. The specificity for these changes was high, ranging from 0.78 to 0.94. Unfortunately, there were insufficient data to calculate sensitivity. This would seem like a novel and exciting means of seizure prediction, but obtaining long-term, reliable recordings of single units is technically challenging. Furthermore, interictal epileptiform activity (spikes) and changes in sleep-wake state will also have a considerable effect on neuronal firing patterns[15], reducing the specificity of such data. It is nevertheless remarkable that there are changes in neuronal firing that are occurring minutes before seizure onset in an area that is distinct from the seizure focus.

This study demonstrates, more than anything, that the careful investigation of affected individuals can give us unprecedented insights into seizure dynamics. With the growing use of single-unit neuronal recordings in humans, we can hope that this will be one in a long line of such studies.

COMPETING FINANCIAL INTERESTS
The author declares no competing financial interests.

1. Jackson, J.H. *West Riding Lunatic Asylum Med. Rep.* **3**, 315–339 (1873).
2. Lennox, W. *Brain* **59**, 113–121 (1936).
3. Matsumoto, H. & Ajmone Marsan, C. *Exp. Neurol.* **9**, 305–326 (1964).
4. Sawa, M., Nakamura, K. & Naito, H. *Electroencephalogr. Clin. Neurophysiol.* **9**, 305–326 (1968).
5. Ishijima, B. *Epilepsia* **13**, 561–581 (1972).
6. Babb, T.L., Wilson, C.L. & Isokawa-Akesson, M. *Electroencephalogr. Clin. Neurophysiol.* **66**, 467–482 (1987).
7. Truccolo, W. *et al. Nat. Neurosci.* **14**, 635–641 (2011).
8. Sypert, G.W., Oakley, J. & Ward, A.A. Jr. *Exp. Neurol.* **28**, 308–325 (1970).
9. Bower, M.R. & Buckmaster, P.S. *J. Neurophysiol.* **99**, 2431–2442 (2008).
10. Bragin, A., Penttonen, M. & Buzsáki, G. *J. Neurosci.* **17**, 2567–2579 (1997).
11. Jasper, H. *Epilepsia* **2**, 91–99 (1961).
12. Lado, F.A. & Moshe, S.L. *Epilepsia* **49**, 1651–1664 (2008).
13. Schindler, K., Elger, C.E. & Lehnertz, K. *Clin. Neurophysiol.* **118**, 1955–1968 (2007).
14. Mormann, F. *et al. Brain.* **130**, 314–33 (2007).
15. Evarts, E.V. *J. Neurophysiol.* **27**, 152–171 (1964).

Calcium channels put synapses in their place

Justin R Fallon

Synapse density and patterning must be tightly regulated to ensure proper circuit formation and function. A new report finds that postsynaptic L-type calcium channels control the pattern and differentiation of developing synapses.

Choreographing functional neural circuits requires the coordinated regulation of synapse formation, patterning, density and maturation[1,2]. Circuits are shaped by the interaction of the intrinsic properties of the synaptic partners, the exchange of biochemical information and electrical activity. Unraveling the underpinnings of circuit formation requires delineation of the precise role of each of these factors. A report by Chen and colleagues in this issue uncovers a surprising new role in synaptic patterning for a familiar component: the dihydropyridine receptor (DHPR) L-type calcium channel[3].

Synaptic patterning, development and maturation is best-understood at the neuromuscular junction[2,4]. Here, entire axon arbors and their postsynaptic partners can be imaged and genetically manipulated over the complete arc of circuit formation. From the earliest times in development, nerve-muscle synapses are arrayed along discrete regions of the muscle in reproducible patterns. Up until 2001, the

The author is in the Department of Neuroscience, Brown University, Providence, Rhode Island, USA.
e-mail: justin_fallon@brown.edu

nature neuroscience

Single-neuron dynamics in human focal epilepsy

Wilson Truccolo[1–4,16], Jacob A Donoghue[1,16], Leigh R Hochberg[1,3–5], Emad N Eskandar[6,7], Joseph R Madsen[8,9], William S Anderson[9], Emery N Brown[10–12], Eric Halgren[13–15] & Sydney S Cash[1]

Epileptic seizures are traditionally characterized as the ultimate expression of monolithic, hypersynchronous neuronal activity arising from unbalanced runaway excitation. Here we report the first examination of spike train patterns in large ensembles of single neurons during seizures in persons with epilepsy. Contrary to the traditional view, neuronal spiking activity during seizure initiation and spread was highly heterogeneous, not hypersynchronous, suggesting complex interactions among different neuronal groups even at the spatial scale of small cortical patches. In contrast to earlier stages, seizure termination is a nearly homogenous phenomenon followed by an almost complete cessation of spiking across recorded neuronal ensembles. Notably, even neurons outside the region of seizure onset showed significant changes in activity minutes before the seizure. These findings suggest a revision of current thinking about seizure mechanisms and point to the possibility of seizure prevention based on spiking activity in neocortical neurons.

Seizures and epilepsy have been recognized since antiquity, yet we continue to struggle to define and understand these paroxysms of neuronal activity. Epileptic seizures are commonly considered to be the result of monolithic, hypersynchronous activity arising from an imbalance between excitation and inhibition in large populations of cortical neurons[1–3]. This view of ictal activity is highly simplified, and the level at which it breaks down is unclear. It is largely based on electroencephalogram (EEG) recordings, which reflect the averaged activity of millions of neurons. Whereas some *in vitro* studies have shown that sparse and asynchronous neuronal activity evolves into a single hypersynchronous cluster with elevated spiking rates at seizure initiation[4,5], as the canonical view would suggest, other animal model studies have supported a much less homogeneous progression in neuronal activity during seizures[6–8]. How well these animal models capture mechanisms operating in human epilepsy remains an open question[9,10]. Very few human studies have gone beyond macroscopic scalp and intracranial EEG signals to examine neuronal spiking underlying seizures[11–14]. Hence the behavior of single neurons in human epilepsy remains largely unknown.

The reliance on the macroscopic information of the EEG has also dominated attempts at discovering physiological changes preceding the seizure. The obvious goal of this approach is in predicting and then preventing seizures[15,16]. While *in vitro* and *in vivo* animal models suggest that different neuronal populations might have distinct roles during a preictal period[4–8,17–21], reliable seizure prediction based on detection of preictal changes in human scalp and intracranial EEG has remained elusive[16]. In addition, most seizures end abruptly and spontaneously, followed by a post-seizure attenuation in EEG activity[22]. The underlying mechanisms governing this behavior are also not understood. Various potential mechanisms, including among others, glutamate depletion, profound inhibition, modulatory effects from subcortical structures and depolarization block, have been hypothesized to underlie seizure termination[22,23]. Although these mechanisms clearly operate at the level of individual cells, to our knowledge, single-unit activity during this period has not been examined in humans. Such information could be useful in developing better strategies for seizure control and preventing status epilepticus[24].

A deeper understanding of neuronal spiking during the different phases of seizure generation would have profound implications for seizure prediction and may provide the basis for new and more effective therapies for people with epilepsy[25]. Here we studied the spiking activity of hundreds of neurons in four persons with focal epilepsy. We found significant changes in preictal activity in subsets of neurons. During seizure initiation and spread, we observed a high degree of heterogeneity in spiking activity. This heterogeneity did not seem to result purely from differences between interneurons and pyramidal cells; heterogeneity was present even within a class. Spiking evolved into a more homogeneous activity across the recorded neuronal ensemble toward seizure termination, during which we observed an almost complete cessation of spiking across the recorded cortical patch. Further, in our data, depolarization block did not seem to have a primary local role during the end of the seizure.

[1]Department of Neurology, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [2]Department of Neuroscience, Brown University, Providence, Rhode Island, USA. [3]Institute for Brain Science, Brown University, Providence, Rhode Island, USA. [4]Rehabilitation Research and Development Service, Department of Veterans Affairs, Providence, Rhode Island, USA. [5]School of Engineering, Brown University, Providence, Rhode Island, USA. [6]Department of Neurosurgery, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [7]Nayef Al-Rodhan Laboratories for Cellular Neurosurgery and Neurosurgical Technology, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [8]Department of Neurosurgery, Children's Hospital and Harvard Medical School, Boston, Massachusetts, USA. [9]Department of Neurosurgery, Brigham and Women's Hospital and Harvard Medical School, Boston, Massachusetts, USA. [10]Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [11]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. [12]Harvard-Massachusetts Institute of Technology, Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. [13]Department of Radiology, University of California, San Diego, San Diego, California, USA. [14]Department of Neurosciences, University of California, San Diego, San Diego, California, USA. [15]Department of Psychiatry, University of California, San Diego, San Diego, California, USA. [16]These authors contributed equally to this work. Correspondence should be addressed to W.T. (wilson_truccolo@brown.edu).

Received 23 December 2010; accepted 15 February 2011; published online 27 March 2011; doi:10.1038/nn.2782

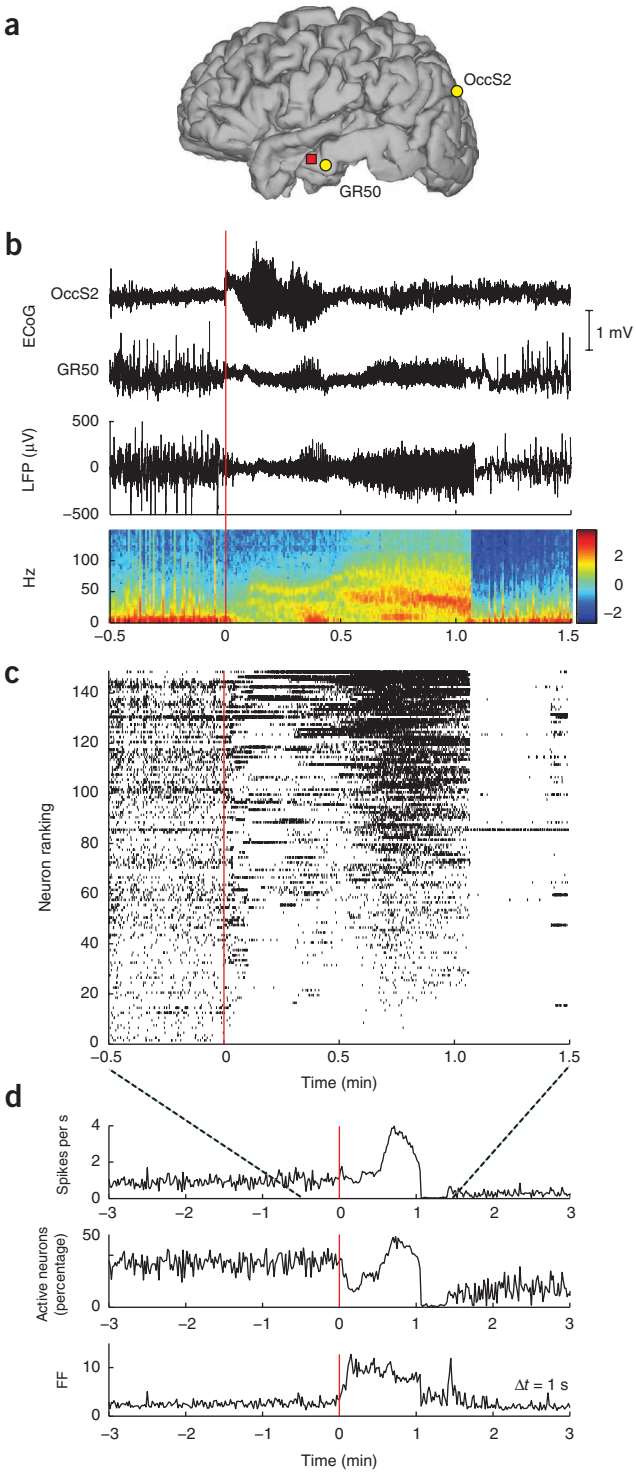

Figure 1 Heterogeneous neuronal spiking patterns during seizure. (**a**) Locations of the microelectrode array in participant A (red square), and subdural ECoG electrodes OccS2 and GR50 in occipital and middle temporal cortices, respectively. (**b**) ECoG traces recorded at the locations shown in **a** during seizure 1. The ECoG-based onset area was identified to be under the occipital electrode OccS2. Seizure onset is at time 0. The local field potential (LFP) recorded from a single channel in the microelectrode array and the corresponding spectrogram (in dB) are shown below. (**c**) Neuronal spike raster plot including all recorded neurons (n = 149). Each hash mark represents the occurrence of an action potential. Neurons were ranked (vertical axis) in increasing order according of their mean spiking rate during the seizure. (This ranking number is unrelated to physical location.) Toward the end of the seizure, activity across the population became more homogeneous until spiking was abruptly interrupted at seizure termination. With the exception of a few neurons, spiking in the recorded population remained suppressed for about 20 s. (**d**) The mean population rate, the percentage of active neurons and the Fano factor (FF) of the spike counts across different neurons at a given time (determined in 1-s time bins). These were roughly stationary during the several minutes preceding the seizure onset. An increase in the Fano factor, reflecting the heterogeneity in neuronal spiking, is observed around seizure onset and precedes an increase in the mean population rate.

RESULTS

We used specialized 96-channel microelectrode arrays[26–31] to record single-unit spiking activity and local field potentials from a 4 mm × 4 mm region of neocortex in four patients with epilepsy refractory to medical treatments. These patients were implanted with subdural grid electrodes to evaluate their cortical EEG activity (electrocorticogram; ECoG) and help localize the onset zone of their seizures for subsequent surgical resection. For research purposes, the microelectrode arrays were placed in addition to the grids (**Fig. 1**). We identified a total of 712 single-unit recordings in four participants (A, B, C and D). Single units were sorted using standard techniques (Online Methods). Although we recorded continuously over several days, the consistent sorting of single units over time periods longer than a few hours proved challenging. Over such long periods, waveforms of extracellularly recorded action potentials could change and units appear or disappear from recordings, owing perhaps to array micromotion, changes in brain states and other factors[32]. For this reason, for each analyzed session we identified for study those single units that were consistently recorded during a time period of ~2 h centered at the onset of each of eight seizures (see Online Methods). Microelectrode arrays in participants A (three seizures; n = 149, 131 and 131 single units recorded), B (two seizures; n = 57 in each) and D (n = 35 in one seizure) were placed in the middle temporal gyrus ~2 cm from the anterior tip. The microelectrode array in participant C (two seizures; n = 82 and 70) was placed in the middle frontal gyrus. Microelectrode arrays were positioned both within (participant C) and outside (~2–4 cm away; participants A, B and D) the seizure onset zone as subsequently defined by the ECoG-electrode locations at which the seizure first appeared (see Online Methods).

Heterogeneous neuronal spiking activity during seizures

Visual inspection of neuronal spike rasters revealed a variety of spiking patterns during seizure initiation, spread and maintenance (**Fig. 1c**; see also **Supplementary Fig. 1** and **Supplementary Movie 1** showing spiking rates on the microelectrode array). For example, whereas some neurons increased their spiking rates near the seizure onset, others decreased. These transient spiking rate modulations occurred at different times for different groups of neurons. The Fano factor (variance divided by the mean) of the single neuron spike counts across the population at a given time during the seizure (1-s time bins), which gives an index of spiking heterogeneity in the ensemble, increased substantially after the seizure onset—in some seizures by fivefold (**Fig. 1d**).

This diversity of neuronal modulation patterns was observed to a greater or lesser extent in all seizures and participants studied (**Supplementary Figs. 1–4**). Such heterogeneity in spiking rate modulation patterns directly challenges the canonical characterization of epileptic seizures as a simple, widespread and homogeneous runaway

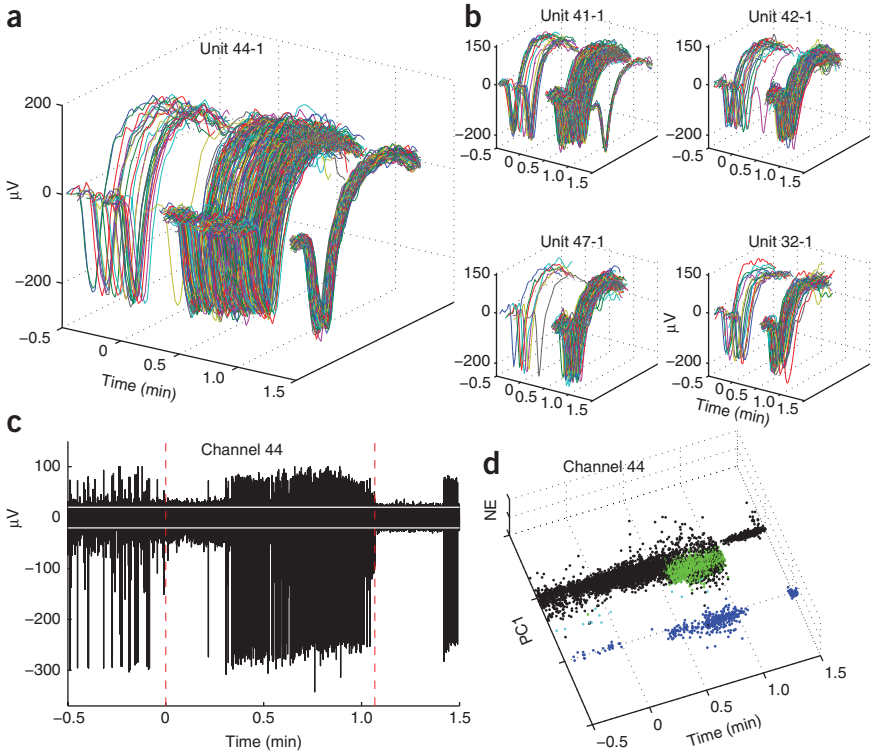

Figure 2 Transient suppression of neuronal spiking during the seizure and at seizure termination. (**a**) Spike waveforms from neuron 44-1 (participant A, seizure 1; neuron ranked no. 131 in **Fig. 1c**). Spiking stopped for ~ 20 s during the initial seizure phase. The lack of major changes in spike waveform and preceding low spiking rate suggest that suppression was not due to sorting artifacts or depolarization block. (**b**) Four examples of units with similar behavior, recorded from different sites. All five units were classified as putative principal cells. (**c**) The high-pass filtered potentials recorded at electrode 44. Larger spikes correspond to unit 44-1, shown in **a**. Dashed vertical lines show seizure onset and termination, respectively. White lines mark ±3 s.d. of the background noise, estimated from the 'silent' period after seizure termination. Another unit with smaller extracellularly recorded action potentials intensifies spiking during the 0.5–1.1 min interval. After seizure termination, both single-unit and multiunit activity were suppressed and the recorded potentials correspond primarily to background noise. Although there is some gradual decrease in spike amplitudes, this decrease is much smaller than what would be expected from depolarization block. See **Supplementary Figure 2** for channels 32, 41, 42 and 47. (**d**) Projection of thresholded waveforms onto a feature space shows clearly separable units. Blue dots represent thresholded spikes from unit 44-1; PC1 and NE denote the first principal component and a nonlinear energy feature, respectively. Green dots correspond to a smaller unit. Black dots correspond to thresholded noise and unsorted multiunit spikes.

excitation leading to a hypersynchronized state. Heterogeneity was present regardless of whether the seizure initiated near (participant C) or far from (the other three participants) the location of the microelectrode array. The admixture of different spiking patterns suggests that heterogeneity is not purely propagation driven but must also reflect local network properties. As based on the Fano factor, this heterogeneity was higher during seizure initiation and decreased toward seizure termination (**Fig. 1d** and **Supplementary Fig. 1**).

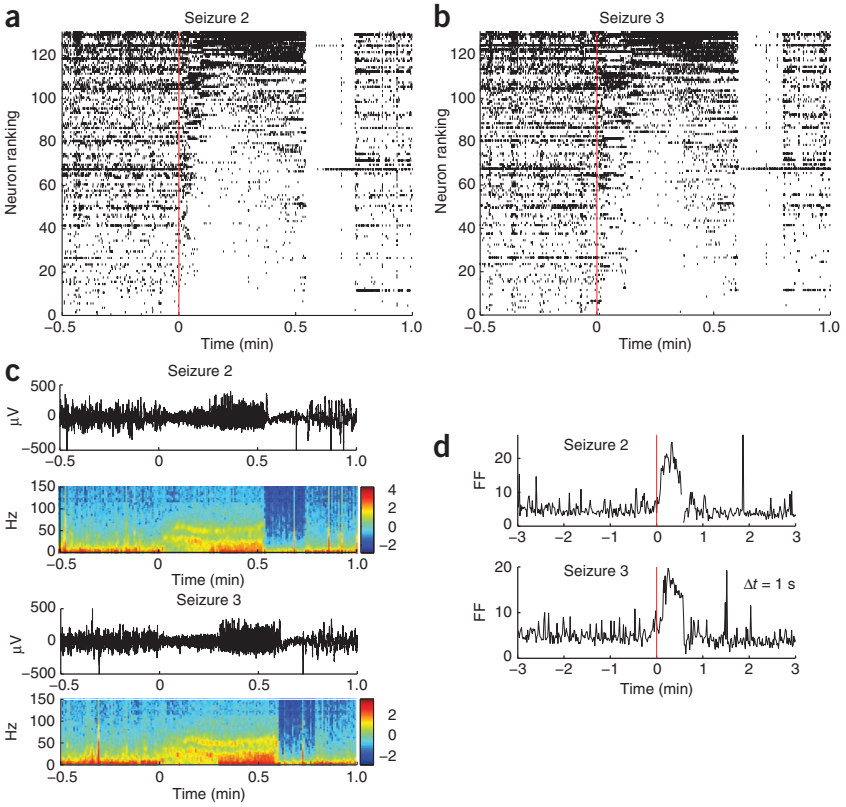

Figure 3 Reproducibility of neuronal spiking modulation patterns across consecutive seizures. (**a**,**b**) An example from participant A with 131 neurons. Following conventions used in **Figure 1c**, neurons are ranked according to their mean rates measured during the seizure. Seizure 3 (**b**) follows the same ranking as seizure 2 (**a**); that is, the single units in any given row of seizures 2 and 3 are the same. Most neurons coarsely preserved the types of spiking rate modulation across the two seizures. For example, the lowest-ranked neurons decreased or stopped spiking; and many of the top-ranked neurons presented similar transient increases in spiking rate modulation. As in seizure 1 (**Fig. 1**), an almost complete suppression of spiking in the neuronal population occurred abruptly at seizure termination. (**c**) The corresponding low-pass filtered local field potentials (LFPs) and spectrograms (from the same microelectrode array channel shown in **Fig. 1**; power in dB). (**d**) The Fano factor for the spike counts (1-s time bins) in the population of recorded neurons showed similar increase during both seizures, reflecting the increased heterogeneity in neuronal spiking across the population.

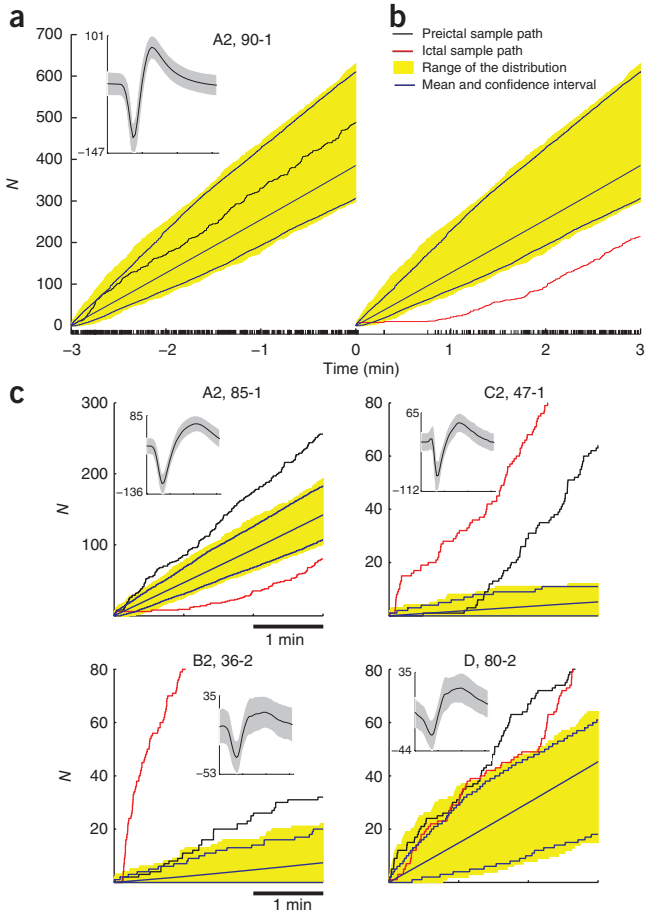

Figure 4 Preictal and ictal modulations in spiking rates. (**a**) The neuronal spiking sample path *N* (neuron 90-1; A2: participant A, seizure 2). The corresponding spike train is shown at the bottom and the inset shows the mean ± 2 s.d. of all recorded spike waveforms. Seizure onset corresponds to time 0. For comparison purposes, the initial value of the sample path is set to 0. The yellow band corresponds to the range of the 3-min-long sample paths observed during a 30-min interictal period preceding the preictal period. Interictal sample paths in this distribution were obtained from an overlapping 3-min-long moving time window, stepped 1 s at a time. Blue curves and surrounding yellow band correspond to the average interictal sample path and the 95% confidence interval, respectively. A sample path was judged to have deviated from the interictal sample paths when it fell outside the range of the collection of interictal sample paths at any given time. (**b**) Neuron 90-1 transiently stopped spiking for tens of seconds just after the seizure onset. As expected, the sample path during the seizure did deviate from the observed interictal paths. The neuron's spiking rate gradually recovered and eventually settled at the typical mean rate. (**c**) Four examples of preictal and ictal sample path deviations, one for each participant. Note that although the preictal and ictal sample paths are plotted along the same axis, they refer to a 3-min period before and after, respectively, the seizure onset.

Seizure termination and suppression of neuronal spiking

In contrast to the beginning of the seizure, seizure termination in participants A, B and C involved widespread, complete cessation of activity for most recorded neurons (**Fig. 1c** and **Supplementary Figs. 1**,**3**). Spiking activity remained suppressed for several seconds (ranging from 5 to 30 s) after seizure termination, until more normal spiking rates gradually returned. Decreases in spiking activity (particularly during either initial or final stages of the seizure) were not due to sorting artifacts such as spike dropout because of obvious alterations in spike waveform. Changes in spike waveforms can be induced by, for example, intense bursting activity. In contrast, **Figure 2a**,**b** shows several examples of units that almost completely ceased spiking for 20–30 s after seizure onset, yet did not present any obvious changes in waveform shape and amplitude that would result in spike dropout. Furthermore, toward the end of the seizure, the same units increased substantially their spiking rates, showing that we were still able to detect their spiking even at much higher spiking rates.

Notably, it also seems that the suppression of these units' spiking either during seizure maintenance or at the end of the seizure was not due to a typical depolarization block scenario[18,19,23] where neuronal spike amplitudes decrease gradually until they cannot be detected or spiking is blocked. For example, the examination of the high-pass-filtered potentials (**Fig. 2c**,**d**) shows that unit 44-1 stopped spiking at a point where peak-to-peak waveform amplitudes were about 300 μV, whereas the threshold for spike detection was set at around −30 μV. (Examples for other neurons are shown in **Supplementary Fig. 2**).

Reproducibility of spike patterns in consecutive seizures

In participants A and B, two consecutive seizures occurred within a period of about an hour. In these cases, we were able to ensure that the same units were recorded and consistently identified during both seizures (see Online Methods). This permitted us to examine the reproducibility of neuronal spiking patterns across seizures. For instance, even though the third seizure in participant A lasted slightly longer than the second, the same motifs in the neuronal spiking patterns recurred (**Fig. 3**). The Pearson correlation coefficient (a measure of similarity) between two spike trains during the initial 30 s of each seizure, for each neuron and averaged across the population, was 0.82. Similar observation was also made for the two consecutive seizures in participant B, with an average correlation coefficient of 0.72 (**Supplementary Fig. 3**), suggesting a high degree of similarity in both seizures. The probability that these correlation coefficients were statistically zero was extremely small ($P < 10^{-6}$; see Online Methods).

We also examined how the similarity between the two ictal events compared to the similarity of random interictal segments preceding the two corresponding consecutive seizures. A random resampling approach using interictal spike rasters each lasting 30 s resulted in significantly smaller correlation coefficients; that is, coefficients whose 95% confidence intervals (mean ± 2 s.d.) corresponded to 0.45 ± 0.14 and 0.16 ± 0.25 for participants A and B, respectively. Despite general similarities across seizures, there were also notable differences in the fraction of active neurons in the population. For example, whereas these fraction increased after a transient decrease at seizure onset in seizure 1 (participant A; **Fig. 1d**), throughout seizures 2 and 3 it remained lower than during the preictal period (**Supplementary Fig. 4**).

Preictal and ictal changes in neuronal spiking activity

To characterize preictal and ictal changes in spiking rates, we represented a single neuron's spike train as a sample path[33]. A sample path consists of the cumulative number of neuronal spikes as a function of time (**Fig. 4a**). The sample path representation allows us to preserve information about transient changes in instantaneous spiking rates. We asked whether an observed sample path for a given neuron during the preictal or ictal periods deviated from the collection of sample paths of same time length observed during a preceding interictal period. In this way, we were able to detect

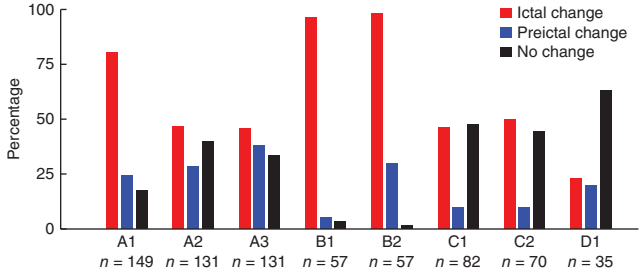

Figure 5 Preictal and ictal sample path deviations with respect to an interictal period. Each bar indicates the percentage of preictal and ictal sample path deviations in the recorded neuronal population, for each participant and seizure. Sample paths and sample path deviations were defined as in **Figure 4**.

transient increases or decreases in spiking rate that were atypical with respect to the preceding interictal activity (**Fig. 4b**,**c**).

We found that substantial numbers of neurons significantly changed their activities as the seizure approached. The percentage of neuronal recordings that showed a preictal deviation varied from 20% (7/35, participant D) to 29.9% (123/411, participant A). The percentages of neurons that showed preictal and/or ictal changes are specified separately for each participant and seizure in **Figure 5**. Observed deviations consisted either of increases or decreases in spiking rate. Across all participants and seizures (712 recordings), 11.8% of recorded neurons increased their spiking rate during the preictal period, and 7.5% decreased. In many cases the onset time for the deviation was earlier than 1 min before the seizure onset time (**Fig. 4c**). This finding suggests that changes in neuronal spiking activity, even for single neurons recorded well outside the identified epileptic focus, may be detected minutes before the seizure onset defined by ECoG inspection. Several neurons in participants A showed consistent sample path deviations across seizures 2 and 3, the two consecutive seizures where the same single units were recorded (**Supplementary Fig. 5**).

The percentage of neuronal recordings that showed a sample path deviation during the ictal period varied from 22.8% (8/35, participant D) to 97.4% (111/114, participant B). In line with the heterogeneity observed in raster plots, several types of deviations were observed during the seizure—all from single units recorded simultaneously from a small cortical patch (encompassing different cortical columns). Across all participants and seizures, 45.4% and 9.9% of the neuronal recordings increased and decreased, respectively, their spiking rates during the seizure. Furthermore, a few neurons showed a transient increase (0.1%) followed by a transient decrease in spiking rates, or vice-versa (1.3%). Overall, we found fifteen different patterns of neuronal modulation when taking into consideration both preictal and ictal patterns. (See **Supplementary Table 1** for detailed percentages for each participant and seizure.) In addition to these types of sample path deviations, a few neurons in participant A also showed much slower and larger modulations in spiking rates that preceded the seizure onset by tens of minutes (**Supplementary Fig. 6**) and were not correlated with obvious behavioral or state changes.

Neurons that showed preictal and ictal modulation in spiking rates tended to have statistically higher bursting rates during interictal periods (Kruskal-Wallis test, $P < 0.01$, Tucker-Kramer correction for multiple comparisons; see Online Methods and **Supplementary Figs. 7** and **8**). Furthermore, in participant A, the single units recorded could be classified into putative interneurons and principal neurons according to spike half-width and peak-to-valley time width features (see Online Methods and **Supplementary Fig. 9**; waveforms from the other three participants formed a homogeneous cluster in each case). On the basis of this classification, 79% (45/57) and 68% (240/354) of interneurons and principal cells recordings, respectively, showed some type of preictal or ictal modulation. In addition, the fraction of recorded interneurons that showed a preictal increase was 60% larger than the corresponding fraction of principal neurons (χ^2 test, $P < 10^{-6}$, with Bonferroni correction for multiple comparisons; see **Supplementary Fig. 10**).

Specificity of preictal changes in neuronal activity

The observed sample path deviations during the preictal period could simply reflect spontaneous or evoked modulations of spiking activity unrelated to the upcoming seizure. Therefore, we also examined how often such sample path deviations could occur during interictal periods. This analysis also provides a preliminary assessment of the specificity of a very simple seizure prediction algorithm that, for example, predicted a seizure every time an observed number of deviations across the recorded neuronal population exceeded a specified threshold. Specifically, we estimated the probability of observing a given number (percentage) of such sample path deviations across the neuronal population during interictal periods. Three different interictal periods were used for participant A, two for participants B and C, and one for participant D. A 3-min-long segment was then randomly selected from a given interictal period and a corresponding (target) sample path was computed for each neuron. Next, we checked whether each sample path deviated from its corresponding distribution derived from paths of same length sampled from a 30-min interictal segment, as done before in **Figure 4**. This 30-min segment was adjacent to but nonoverlapping with the randomly selected target sample path. The entire procedure was repeated hundreds of times to obtain a distribution of the percentage of deviations across the recorded neuronal population during a given interictal period. Finally, a threshold was defined based on the average number of preictal deviations across the examined seizures for a given participant. (For example, the threshold for participant A was derived from the mean of the percentage of neurons that showed a preictal deviation across the three seizures; **Fig. 5**.) Given this threshold and a distribution of percentage of deviations, we could then compute the probability that a given percentage of observed deviations across the population during an interictal period was smaller than the specified threshold. As mentioned above, this probability provides an estimate of the specificity of a seizure prediction algorithm based on the defined sample path deviations.

With the exception of participant C, for whom interictal activity showed a very high rate of epileptiform events consisting of slow neuronal bursting and ECoG spike and wave discharges, promising specificity results (0.78–0.94) were obtained for participants A, B and D (**Supplementary Fig. 11**). Despite these positive preliminary results, we emphasize that a more conclusive assessment will require much larger data sets, to probe a larger range of physiological and behavioral states, as well as corresponding sensitivity analyses.

DISCUSSION

Our findings, based on the most extensive description of single-unit activity in human neocortical seizures yet reported, reveal several important and heterodox points about the nature of epileptic activity. First, the observed heterogeneity in neuronal behavior argues against homogeneous runaway excitation or widespread paroxysmal depolarization as the primary mechanism underlying seizure initiation. Rather, our data indicate that seizures result from a complex interplay among groups of neurons that present different types of spiking

behaviors evolving at multiple temporal and spatial scales. We have also observed similar heterogeneity in interictal discharges[31]. Given the 1-mm microelectrode length, it is likely we recorded from cells in layers 3 and 4. Several studies[5,34,35] suggest that epileptiform activity involves and is perhaps initiated by cells in layer 5. Although the potential role of these cells needs to be further explored, they do not seem to be driving homogenous activity. Furthermore, because of the nature of these recordings, we might not have recorded from the initiation site in any of the participants (including participant C). As a result, it is possible that a different form of neuronal dynamics would be observed at the 'focus'. This is especially likely to be the case in a region of dysplasia in which the neurons and their layering are severely abnormal.

Heterogeneity in neuronal spiking activity during seizures has been previously observed in animal model studies[6–8]. In particular, it has been hypothesized[7] that such heterogeneity could reflect three main sequential stages or states during seizure spreading: 'depressed', 'projected' and 'propagated' states. The fact that the heterogeneity reported here in human epilepsy appears simultaneously in small patches of neocortex would speak against this hypothesis; or, alternatively, it would require that different groups of neurons entered and dwelled in different states with different time constants, perhaps owing to differences in initial conditions or in intrinsic and local network dynamical properties. In addition, our finding based on the classification of units into putative principal cells and interneurons, namely that some putative pyramidal cells increased while other decreased their spiking rates, suggests that such heterogeneity does not simply reflect interleaved spiking of pyramidal cells and interneurons[18,19]. We also speculate that the observed heterogeneity in the neuronal collective dynamics[30] could result from fragmentation into multicluster synchronization, which has been studied in various dynamical systems[36]. The fact that such diverse spiking activity underlies seemingly 'monomorphic' EEG waveforms raises the possibility that even normal cortical rhythms might also reflect very heterogeneous underlying neuronal ensemble spiking.

Second, one of the noteworthy features of these data was the abrupt and widespread suppression of neuronal action potentials at seizure end. In participant A, for example, spiking of both putative interneurons and principal neurons became suppressed. Previous work has suggested that seizures might self-terminate because of depolarization block resulting from changes in ionic concentrations in the extracellular space. A depolarization block of neuronal spikes due to a chain of events that ends with astrocytic release of large amounts of potassium has been hypothesized[23]. However, several units became suppressed at seizure termination without showing typical signatures of depolarization block (**Fig. 2**)—an indication that depolarization block was not the primary local factor responsible for the observed marked suppression in spiking activity. Furthermore, our results also argue against massive inhibition from a local source because the suppression of neuronal spiking affected both putative interneurons and other principal cells. Alternatively, distant modulatory inputs involving subcortical structures—for example, thalamic nuclei or substantia nigra pars reticulata[22]—could lead to seizure termination. There is also the possibility that this critical transition could arise from an emergent property of the large-scale network itself leading to spatially synchronous extinction[37].

From a therapeutic perspective, our analysis demonstrates, for the first time in humans, that preictal neuronal spiking reflects a distinct and widely occurring physiological state in focal epilepsies. This is true even outside the region of seizure initiation, suggesting that it may be possible to obtain predictive information from individual neuronal activity without necessarily localizing what has been traditionally considered the seizure focus. Substantiation of this possibility will require large data sets, perhaps only available through multisite collaborative efforts, containing sufficient interictal data for proper specificity and sensitivity analyses[16]. Promisingly, our data suggests that the clinical care of patients with epilepsy could be revolutionized by using dynamics of ensembles of single neurons to predict seizures.

METHODS

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

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

The authors thank the patients who participated in this study, as well as the nursing and physician staff at each facility. We also thank A.M. Chan, C.J. Keller, A. Dykstra and J.E. Cormier for technical assistance, and J.P. Donoghue and K.J. Staley for critical reading of the manuscript. This research is funded by a CIMIT grant and US National Institutes of Health (NIH) National Institute of Neurological Disorders and Stroke (NINDS) NS062092 to S.S.C.; an NIH–NINDS Career Award (5K01NS057389) to W.T.; NIH NS018741 to E.H.; NINDS K08NS066099-01A1 to W.S.A.; US National Eye Institute EY017658, US National Institute on Drug Abuse NS063249, US National Science Foundation IOB 0645886, Howard Hughes Medical Institute and the Klingenstein Foundation to E.N.E.; NIH Director's Pioneer Award DP1OD003646 to E.N.B.; US Department of Veterans Affairs Career Development Transition Award, Doris Duke Charitable Foundation–Clinical Scientist Development Award, Massachusetts General Hospital–Deane Institute for Integrated Research on Atrial Fibrillation and Stroke, and NIH-NIDCD R01DC009899 to L.R.H. The contents do not represent the views of the Department of Veterans Affairs or the United States government.

AUTHOR CONTRIBUTIONS

W.T., S.S.C. and J.A.D. wrote the paper. W.T. and J.A.D. conducted the data analysis. Data collection and preprocessing were done by J.A.D., W.T. and S.S.C. S.S.C., L.R.H., W.T. and E.H. conceived and planned the research. E.N.B. provided guidance on methods of data analysis and interpretation. E.N.E., W.S.A. and J.R.M. performed the surgeries and microelectrode array implantations. All authors participated in editing the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare competing financial interests: details accompany the full-text HTML version of the paper at http://www.nature.com/natureneuroscience/.

Published online at http://www.nature.com/natureneuroscience/.
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions/.

1. Penfield, W.G. & Jasper, H.H. *Epilepsy and the Functional Anatomy of the Human Brain* (Little, Brown, Boston, 1954).
2. Schwartzkroin, P.A. Basic mechanisms of epileptogenesis. in *The Treatment of Epilepsy* (ed. Wyllie, E.) 83–98 (Lea and Febiger, Philadelphia, 1993).
3. Fisher, R.S. *et al.* Epileptic seizures and epilepsy: definitions proposed by the International League Against Epilepsy (ILAE) and the International Bureau for Epilepsy (IBE). *Epilepsia* **46**, 470–472 (2005).
4. Jiruska, P. *et al.* High-frequency network activity, global increase in neuronal activity, and synchrony expansion precede epileptic seizures *in vitro*. *J. Neurosci.* **30**, 5690–5701 (2010).
5. Pinto, D.J., Patrick, S.L., Huang, W.C. & Connors, B.W. Initiation, propagation, and termination of epileptiform activity in rodent neocortex *in vitro* involve distinct mechanisms. *J. Neurosci.* **25**, 8131–8140 (2005).
6. Matsumoto, H. & Ajmone Marsan, C. Cortical cellular phenomena in experimental epilepsy: ictal manifestations. *Exp. Neurol.* **9**, 305–326 (1964).
7. Sawa, M., Nakamura, K. & Naito, H. Intracellular phenomena and spread of epileptic seizure discharges. *Electroencephalogr. Clin. Neurophysiol.* **24**, 146–154 (1968).
8. Bower, M. & Buckmaster, P.S. Changes in granule cell firing rates precede locally recorded spontaneous seizures by minutes in an animal model of temporal lobe epilepsy. *J. Neurophysiol.* **99**, 2431–2442 (2008).
9. Jefferys, J.G.R. Models and mechanisms of experimental epilepsies. *Epilepsia* **44** (suppl. 12): 44–50 (2003).
10. Buckmaster, P.S. Laboratory animal models of temporal lobe epilepsy. *Comp. Med.* **54**, 473–485 (2004).

11. Halgren, E., Babb, T.L. & Crandall, P.H. Post-EEG seizure depression of human limbic neurons is not determined by their response to probable hypoxia. *Epilepsia* **18**, 89–93 (1977).
12. Wyler, A.R., Ojemann, G.A. & Ward, A.A. Jr. Neurons in human epileptic cortex: correlation between unit and EEG activity. *Ann. Neurol.* **11**, 301–308 (1982).
13. Babb, T.L., Wilson, C.L. & Isokawa-Akesson, M. Firing patterns of human limbic neurons during stereoencephalography (SEEG) and clinical temporal lobe seizures. *Electroencephalogr. Clin. Neurophysiol.* **66**, 467–482 (1987).
14. Engel, A.K., Moll, C.K.E., Fried, I. & Ojeman, G.A. Invasive recordings from the human brain: clinical insights and beyond. *Nat. Rev. Neurosci.* **6**, 35–47 (2005).
15. Lopes da Silva, F.H. *et al.* Dynamical diseases of brain systems: different routes to epilepsy. *IEEE Trans. Biomed. Eng.* **50**, 540–548 (2003).
16. Mormann, F., Andrzejak, R.G., Elger, C.E. & Lehnertz, K. Seizure prediction: the long and winding road. *Brain* **130**, 314–333 (2007).
17. Bragin, A., Engel, J. Jr., Wilson, C.L., Fried, I. & Mathern, G.W. Hippocampal and entorhinal cortex high-frequency oscillations (100–500 Hz) in human epileptic brain and in kainic acid-treated rats with chronic seizures. *Epilepsia* **40**, 127–137 (1999).
18. Bikson, M., Hahn, P.J., Fox, J.E. & Jefferys, J.G.R. Depolarization block of neurons during maintenance of electrographic seizures. *J. Neurophysiol.* **90**, 2402–2408 (2003).
19. Ziburkus, J., Cressman, J.R., Barreto, E. & Schiff, S.J. Interneuron and pyramidal cell interplay during in vitro seizure-like events. *J. Neurophysiol.* **95**, 3948–3954 (2006).
20. Cymerblit-Sabba, A. & Schiller, Y. Network dynamics during development of pharmacologically induced epileptic seizures in rats in vivo. *J. Neurosci.* **30**, 1619–1630 (2010).
21. Yaari, Y. & Beck, H. "Epileptic neurons" in temporal lobe epilepsy. *Brain Pathol.* **12**, 234–239 (2002).
22. Lado, F.A. & Moshe, S.L. How do seizures stop? *Epilepsia* **49**, 1651–1664 (2008).
23. Bragin, A., Penttonen, M. & Buzsaki, G. Termination of epileptic afterdischarge in the hippocampus. *J. Neurosci.* **17**, 2567–2579 (1997).
24. Treiman, D.M. Status epilepticus. in *The Treatment of Epilepsy: Principles and Practice.* (ed. Wyllie, E.) 681–697 (Lippincott Williams & Wilkins, Philadelphia, 2001).
25. Jacobs, M.P. *et al.* Curing epilepsy: progress and future directions. *Epilepsy Behav.* **14**, 438–445 (2009).
26. Hochberg, L.R. *et al.* Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* **442**, 164–171 (2006).
27. Truccolo, W., Friehs, G.M., Donoghue, J.P. & Hochberg, L.R. Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *J. Neurosci.* **28**, 1163–1178 (2008).
28. Schevon, C.A. *et al.* Microphysiology of epileptiform activity in human neocortex. *J. Clin. Neurophysiol.* **25**, 321–330 (2008).
29. Kim, S.-P., Simeral, J.D., Hochberg, L.R., Donoghue, J.P. & Black, M.J. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J. Neural Eng.* **5**, 455–476 (2008).
30. Truccolo, W., Hochberg, L.R. & Donoghue, J.P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13**, 105–111 (2010).
31. Keller, C.J. *et al.* Heterogeneous neuronal firing patterns during interictal epileptiform discharges in the human cortex. *Brain* **133**, 1668–1681 (2010).
32. Santhanam, G. *et al.* Hermes B: a continuous neural recording system for freely behaving primates. *IEEE Trans. Biomed. Eng.* **54**, 2037–2050 (2007).
33. 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 (2005).
34. Connors, B.W. Initiation of synchronized neuronal bursting in neocortex. *Nature* **310**, 685–687 (1984).
35. Ulbert, I., Heit, G., Madsen, J., Karmos, G. & Halgren, E. Laminar analysis of human neocortical interictal spike generation and propagation: current source density and multiunit analysis in vivo. *Epilepsia* **45** (suppl. 4): 48–56 (2004).
36. Amritkar, R.E. & Rangarajan, G. Stability of multicluster synchronization. *Int. J. Bifurc. Chaos* **19**, 4263–4271 (2009).
37. Amritkar, R.E. & Rangarajan, G. Spatially synchronous extinction of species under external forcing. *Phys. Rev. Lett.* **96**, 258102 (2006).

ONLINE METHODS

Participants, clinical electrodes and recordings. Four patients (A, B, C and D, ages 21–52 years (mean 29.7), three women) with medically intractable focal epilepsy underwent clinically indicated intracranial cortical recordings using grid electrodes for epilepsy monitoring[38,39]. Clinical electrode implantation, positioning, duration of recordings and medication schedules were based purely on clinical need as judged by an independent team of clinicians. Patients were implanted with intracranial subdural grids, strips and/or depth electrodes (Adtech Medical Instrument Corporation) for 5–10 d in a specialized hospital setting until data sufficiently identified the seizure focus for appropriate resection. Continuous intracranial EEG was recorded with standard recording systems (XLTEK) and captured many seizures. Seizure onset times were determined by an experienced encephalographer (S.S.C.) through inspection of ECoG recordings, referral to the clinical report of the ECoG and clinical manifestations recorded on video. The number of seizures varied across the participants. Owing to operational issues, not all of these seizures were recorded or provided data with a high signal-to-noise ratio. Among data sets with clearly separable single units, we selected eight different seizures among the four participants. All steps of the analysis of intracranial EEG data were performed using Neuroscan Edit software (Compumedics) and custom designed Matlab (MathWorks) software.

Participant A. Participant A was a right-handed woman 52 years old at the time of the continuous 8-d video and invasive EEG monitoring study. She had a history of complex partial seizures with occasional secondary generalization beginning at the age of 4 and typically suffered from 10–15 events per day. Her seizures usually presented with sudden speech arrest associated with confusion and repetition of the activity she was doing just before the onset of the seizure. Magnetic resonance imaging (MRI) showed a large lesion in the left hemisphere extending from the occipital region to the temporal region, which was consistent with encephalomalacia. Moreover, a positron-emission tomography (PET) scan showed hypometabolism in the left occipital, temporal and parietal regions. On the basis of these findings, the patient was implanted with a combination of a subdural grid and strip electrodes and depth electrodes over the left hemisphere. The NeuroPort microelectrode array was placed in the middle temporal gyrus in a region of cortex nearly certain to be resected. The distance to the nearest ECoG electrode where seizure onsets were detected was ~2 cm.

There were several seizures throughout the recordings; all lasted about 1 min or less and had a similar electrographic pattern. They typically began with low-amplitude, fast activity—up to 300 μV, ~30 Hz—in an anterior temporal strip, sometimes with an associated burst of polyspike activity in an occipital strip, or began at an occipital site (**Fig. 1**). Simultaneously, there was a generalized suppression of the grid activity beginning in the posterior inferior quadrant and then spreading to encompass the entire grid. After the onset, there was a buildup of high frequencies and higher amplitudes (up to 500 μV) in the posterior inferior quadrant and occipital region. Sometimes, rhythmic ECoG spikes occurred in the anterior temporal strip and posterior temporal depths.

At the conclusion of the study the patient underwent resection of the left anterior temporal lobe. Pathology showed hippocampal sclerosis with secondary cortical gliosis. The patient remained seizure free for 1 year after the resection, but seizures returned after this period.

Participant B. Participant B was a right-handed man 21 years old at the time of the continuous 8-d video and invasive EEG monitoring study. He had suffered from seizures since at least age 15. Most of his events were characterized by a blank stare and oral automatisms accompanied by stiffening and posturing of the right hand. He was implanted with a series of strip electrodes covering the left frontal and temporal regions. The NeuroPort array was placed in the middle temporal gyrus about 1–2 cm posterior to the temporal tip. The distance to the nearest ECoG electrode where seizure onsets were detected was ~2 cm. Sharp waves were observed interictally in the posterior temporal regions. Four seizures, all with similar clinical and electrographic features, were recorded. Seizures were characterized by a left gaze preference and tonic and then clonic movements of the right arm. Electrographically, the seizures began with a generalized burst of sharp waves followed by sharp wave complexes that were maximal in mesial temporal leads. The patient subsequently underwent a left temporal lobectomy. Pathological examination of the tissue revealed mild dysplastic changes in the lateral temporal neocortex and gliosis and moderate neuronal loss in regions CA4 and CA3 of the hippocampus. Thirteen months after his surgery, he remains seizure free.

Participant C. Participant C was a right-handed woman 22 years old at the time of the continuous 8-d video and invasive EEG monitoring. She had a history of partial seizures with rapid generalization beginning around 14 years of age and typically suffered four to seven attacks per day. Several different clinical manifestations of her seizures had been observed, including events that would occur out of sleep and consist of screaming and whole body shaking. She also had spells of staring, unresponsiveness and oral automatisms. These occasionally progressed to generalized tonic-clonic activity or atonia. Imaging studies including MRI and PET showed diffuse atrophy and questionable abnormalities in the right frontal lobe. Previous EEG monitoring suggested a right-sided region of epileptogenesis, emanating from the either the right frontal or temporal lobe, so she was implanted with extensive subdural coverage of the right hemisphere. The NeuroPort array was implanted in the middle frontal gyrus within a broad onset region.

There were a total of 30 definite seizures throughout the recording session. Each clinical seizure was characterized by a subtle head turn and drop to the right or left that lasted 10–15 s. Electrographically, the seizures fell into four different groups depending on the presence of generalized EEG spikes as well as the seizure onset evolution. The seizures incorporated into our analysis are described in detail below:

Seizure 1 began with an onset of generalized epileptiform spiking activity, followed by 1–2 s of generalized attenuation and then prominent spike and wave activity. Clinically, the subtle head turn corresponded to the generalized EEG spike discharges. The electrographic seizure lasted approximately 11 s.

Seizure 2 started with a generalized EEG spike followed by 1–2 s of attenuation and then the onset of a generalized, high-frequency buzz lasting 4 s. This was followed by EEG spike and wave discharges in the frontopolar strip, the subfrontal strip and segments of the grid of electrodes, lasting 5 s. The seizure corresponded clinically to subtle head movements after the onset of the seizure. The electrographic seizure lasted approximately 11 s.

At the conclusion of the study the patient underwent resection of the left temporal lobe. Pathology confirmed the lesion to be a cortical dysplasia. Six months after surgery the patient had experienced a reduction in the overall number of seizures but continued to have episodes of behavioral arrest.

Participant D. Participant D was a right-handed woman 24 years old at the time of the continuous 4-d video and invasive EEG monitoring study. She had a history of complex partial seizures beginning at the age of fourteen. Her seizures tended to present with staring spells, gulping sounds, a general feeling of heat and hand automatisms. She often had no awareness of the events. Previous EEG monitoring suggested that these seizures originated from the left anterior temporal region. MRI showed a large lesion in the left temporal lobe, believed to be a glioma. In accordance with the above findings, the patient was implanted with a combination of subdural and depth electrodes focused on the left temporal region. The NeuroPort research electrode was placed in the middle temporal gyrus about 2 cm posterior to the temporal tip. The distance to the nearest ECoG electrode where seizure onsets were detected was ~4 cm.

There were five electrographic seizures captured over the course of the 4-d recording session: three typical clinical seizures lasting 20–40 s and two subclinical seizures of shorter duration. The three clinical seizures manifested with staring spells and the participant's report of a general sensation of heat. The EEG revealed the electrographic onset for each of these seizures in the anterior temporal region with a typical subsequent spread across the superior and middle anterior strip electrodes, as well as to the depth electrodes.

At the conclusion of the study, the patient underwent a left anterior temporal lobectomy, a left amygdalohippocampectomy and a lesionectomy of the posterior temporal lobe. Pathology determined the lesion to be a low-grade glioneuronal tumor. At last evaluation, 6 months after surgery, the patient was seizure free.

Microelectrode array location, recordings and analysis. Approval for these studies was granted by local Institutional Review Boards (Partners Human Research Committee) and participants were enrolled after informed consent was obtained. The implanted NeuroPort array (Blackrock Microsystems), which has been used in several previous studies[26–31], is a 4 mm × 4 mm microelectrode array composed of 100 platinum-tipped silicon probes. Arrays with 1.0-mm

doi:10.1038/nn.2782

electrode lengths were implanted in the middle frontal gyrus (participant C) and middle temporal gyrus (participants A, B and D). The array's distance to the nearest ECoG electrode containing seizure onsets, calculated using MRI and post-operative computed tomography registration, was calculated as ~2 cm in participants A and B, ~4 cm in participant D, and within a broad seizure onset zone for participant C. Histology after resection confirmed that the tips of the electrodes were in the lower portion of layer III in two of the three cases (A and B). In these two cases, the histology around the array appeared normal. In a third participant (D), the histology of the cortex also appeared normal but the exact path of the electrode could not be reconstructed. In the fourth participant (C), the array was placed in an area of cortex that was the suspected focus and looked abnormal on visual inspection during the time of the initial craniotomy. Ultimately, however, this area was not resected. For detection and extraction of extracellularly recorded action potentials, the analog signal (0.3 Hz–7.5 kHz) from each of the 96 active electrodes was sampled at 30 kHz, online digitally high-pass filtered (250 Hz–7.5 kHz, fourth-order Butterworth filter) and automatically amplitude thresholded (Cerebus system, Blackrock Microsystems). The extracted waveforms were sorted using standard methods[40] and Offline Sorter (Plexon) and were tracked to ensure that spike rate changes were not obvious artifacts due to waveform changes. Spike sorting was performed on projections of waveforms into feature spaces given by principal components and a nonlinear energy measure. (This nonlinear energy measure is akin to the product of the raw amplitude times the derivative of the signal; Offline Sorter User Guide, version 3, Plexon). In the case of consecutive seizures occurring within ~1 h, we were able to track the neuronal spike waveforms and ensure that the same unit was being isolated. Single-neuron bursting rates were computed as previously described[41]. Classification of neurons into putative interneurons or principal cells[42] was based on neuronal spike waveforms extracted after offline high-pass filtering (forward-backward Kaiser filter, 250 Hz high-pass cutoff) of the original data. Local field potentials were low-pass filtered at 250 Hz. The probability that Pearson correlation coefficients were statistically different from 0 was computed using the Matlab function corrcoef.m (MathWorks).

38. Delgado-Escueta, A.V. & Walsh, G.O. The selection process for surgery of intractable complex partial seizures: surface EEG and depth electrography. in *Epilepsy* (eds. Ward, A.A. Jr., Penry, J.K. & Purpura, D.P.) 295–326 (Raven, New York, 1983).
39. Engel, J., Crandall, P.H. & Rausch, P. Surgical treatment of partial epilepsies. in *The Clinical Neurosciences* (eds. Rosenburg, R.N., Grossman, R.G. & Schoclet, S.) 1349–1380 (Churchill Livingston, New York, 1983).
40. Lewicki, M.S. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network* **9**, 53–78 (1998).
41. Staba, R.J., Wilson, C.L., Bragin, A., Fried, I. & Engel, J. Jr. Sleep states differentiate single neuron activity recorded from human epileptic hippocampus, entorhinal cortex, and subiculum. *J. Neurosci.* **22**, 5694–5704 (2002).
42. Barthó, P. *et al.* Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J. Neurophysiol.* **92**, 600–608 (2004).

Single-neuron dynamics in human focal epilepsy

SUPPLEMENTARY INFORMATION

Wilson Truccolo[1–4,16], Jacob A Donoghue[1,16], Leigh R Hochberg[1,3-5], Emad N Eskandar[6,7], Joseph R Madsen[8,9], William S Anderson[9], Emery N Brown[10–12], Eric Halgren[13–15] & Sydney S Cash[1]

[1]Department of Neurology, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [2]Department of Neuroscience, Brown University, Providence, Rhode Island, USA. [3]Institute for Brain Science, Brown University, Providence, Rhode Island, USA. [4]Rehabilitation Research and Development Service, Department of Veterans Affairs, Providence, Rhode Island, USA. [5]School of Engineering, Brown University, Providence, Rhode Island, USA. [6]Department of Neurosurgery, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [7]Nayef Al-Rodhan Laboratories for Cellular Neurosurgery and Neurosurgical Technology, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [8]Department of Neurosurgery, Children's Hospital and Harvard Medical School, Boston, Massachusetts, USA. [9]Department of Neurosurgery, Brigham and Women's Hospital and Harvard Medical School, Boston, Massachusetts, USA. [10]Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts, USA. [11]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. [12]Harvard-Massachusetts Institute of Technology, Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. [13]Department of Radiology, University of California, San Diego, San Diego, California, USA. [14]Department of Neurosciences, University of California, San Diego, San Diego, California, USA. [15]Department of Psychiatry, University of California, San Diego, San Diego, California, USA. [16]These authors contributed equally. Corresponding author: Wilson Truccolo (wilson_truccolo@brown.edu)

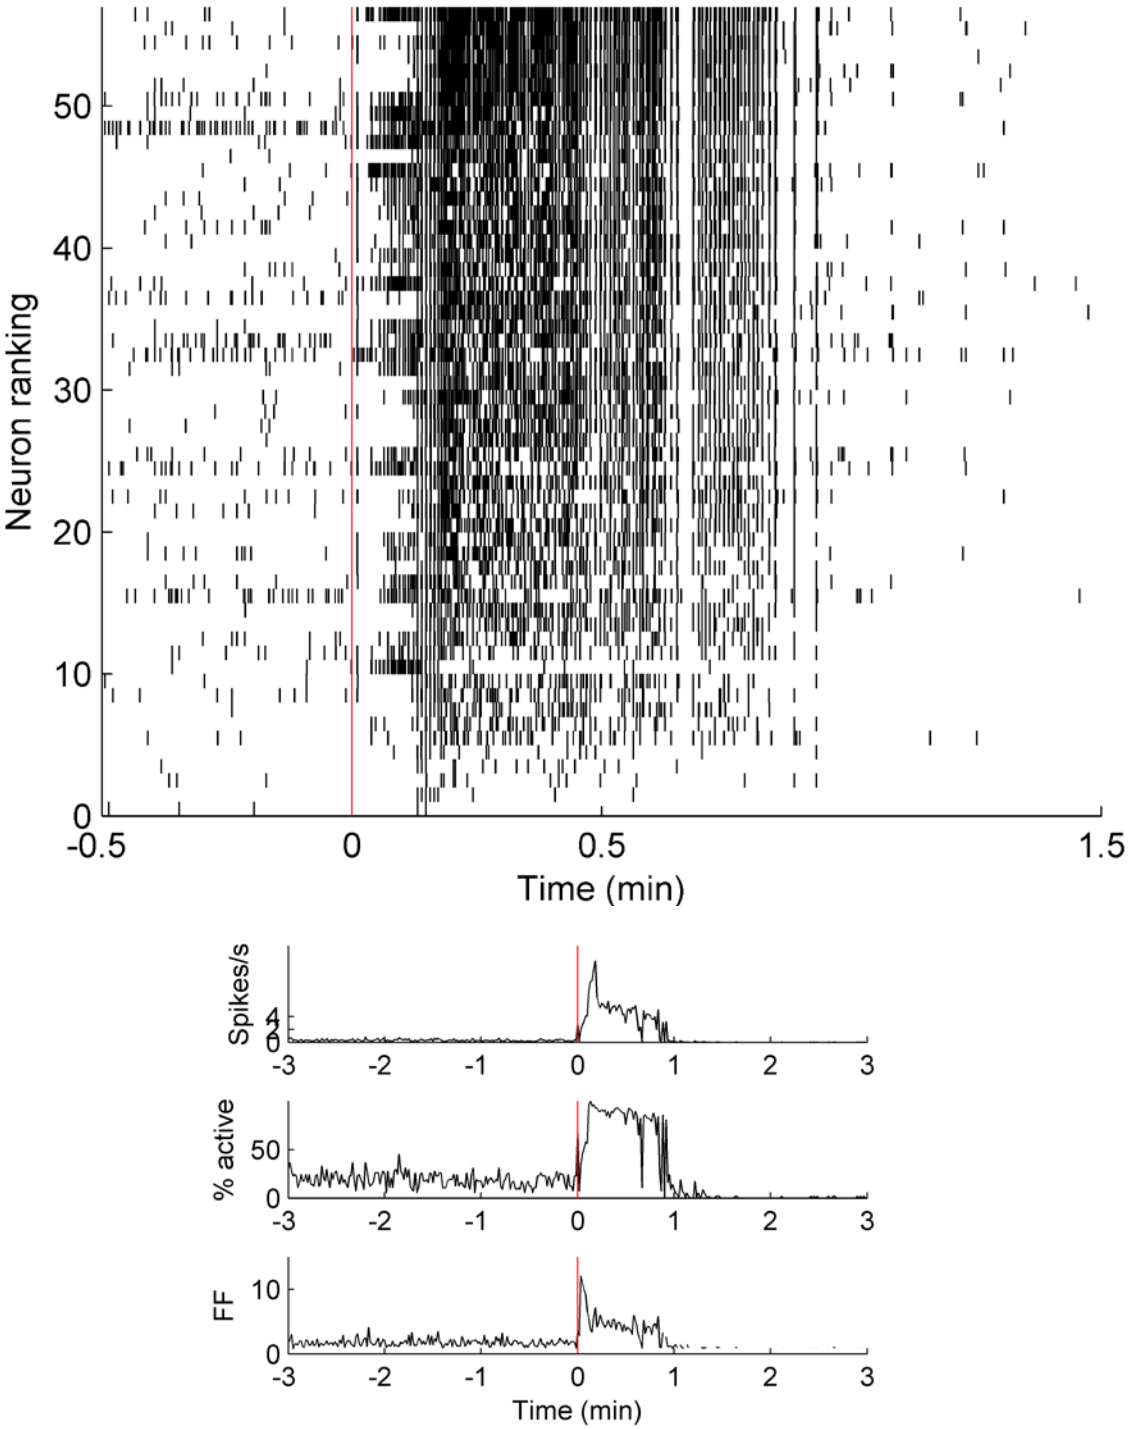

Supplementary Figure 1. Heterogeneous spiking patterns during seizure evolution (additional examples). Heterogeneity of spiking activity in Patient B, seizure 1, is higher during the early stages of seizure initiation, as also reflected in the peak in the Fano factor (FF) of the spike counts (in 1-second time bins) across the population. The population spiking gradually increases during this period as well as the percentage of active neurons. (Percentage of active neurons was computed on a 1-second time scale, i.e. a neuron was considered active once if spiked at least once in a given 1-second time bin.) At the end of this initial period, about 60% of the neurons had become active and this percentage jumped to the activation peak (~ 97%). Nevertheless, even at this time, we can still observe heterogeneity in spiking activity. For example, the neuron ranked 11 essentially 'shuts down' after this peak activity.

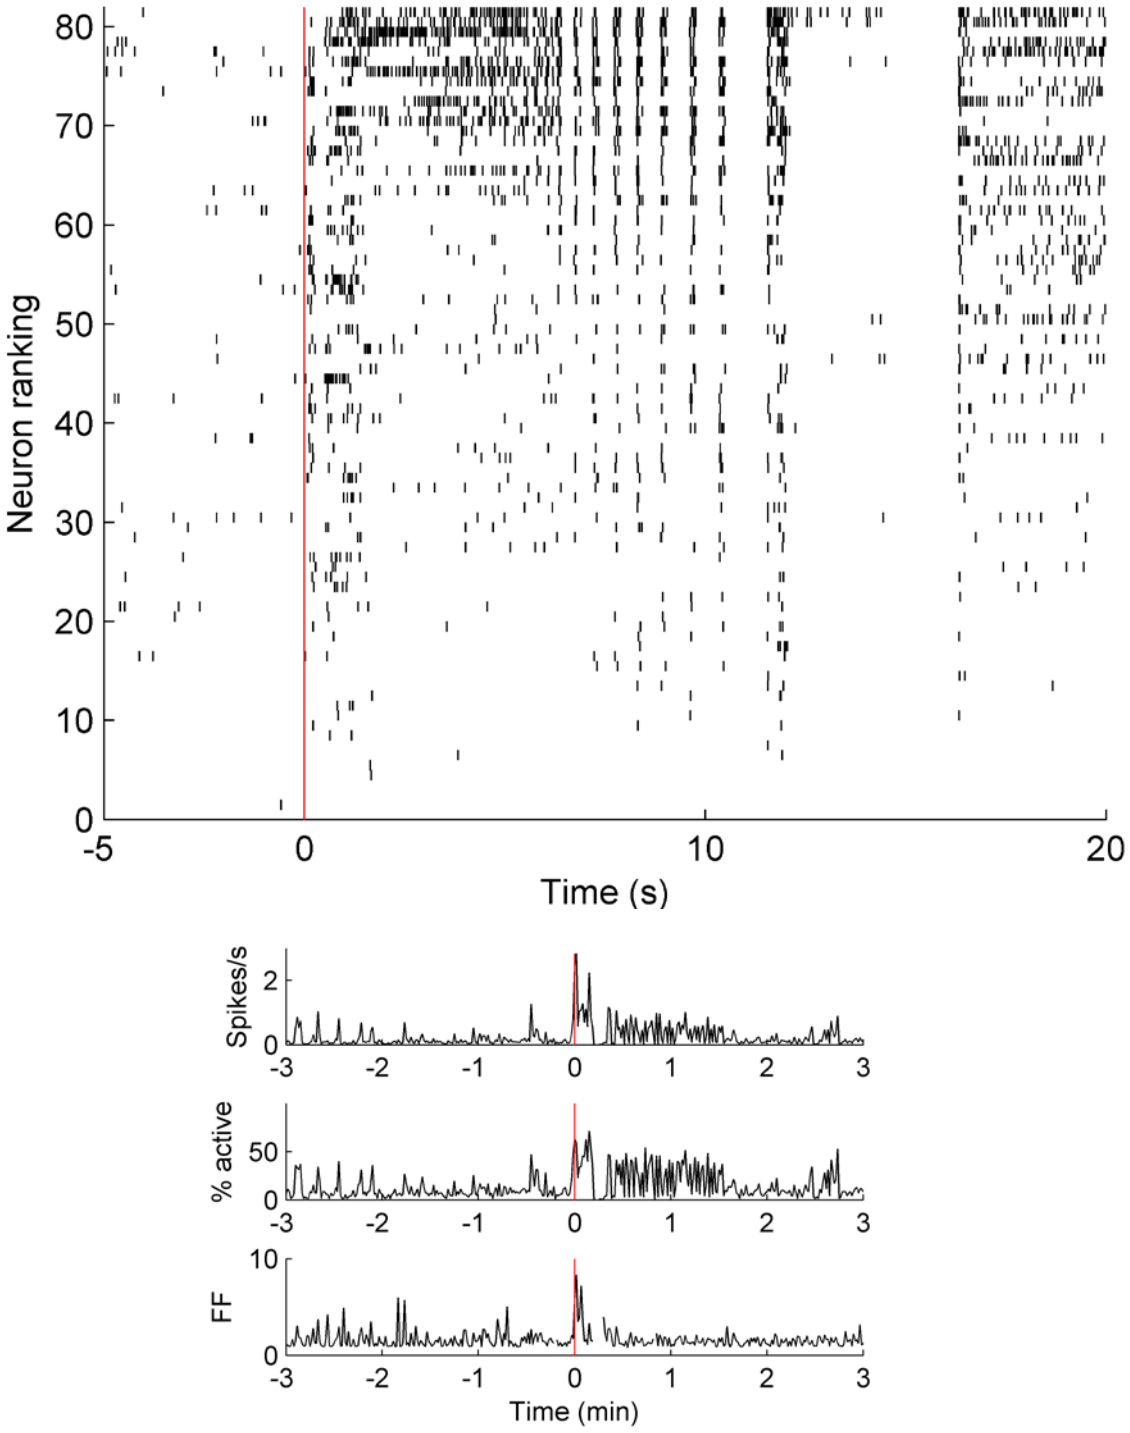

Supplementary Figure 1 (Cont.). Seizure 1 in Patient C had a relatively short duration (~ 11 s). Heterogeneous spiking behavior is most prominent during the first 5 seconds of the seizure. The Fano factor peaks near seizure onset and shows also several transient increases preceding the seizure. During the second half of the seizure, several synchronized bursts of activity can also be seen in the population spike rate and in the percentage of active neurons. These bursts, interspaced with brief silences, resemble failed seizure terminations. After a postictal silence lasting ~ 5 s, a brief period of higher activity follows.

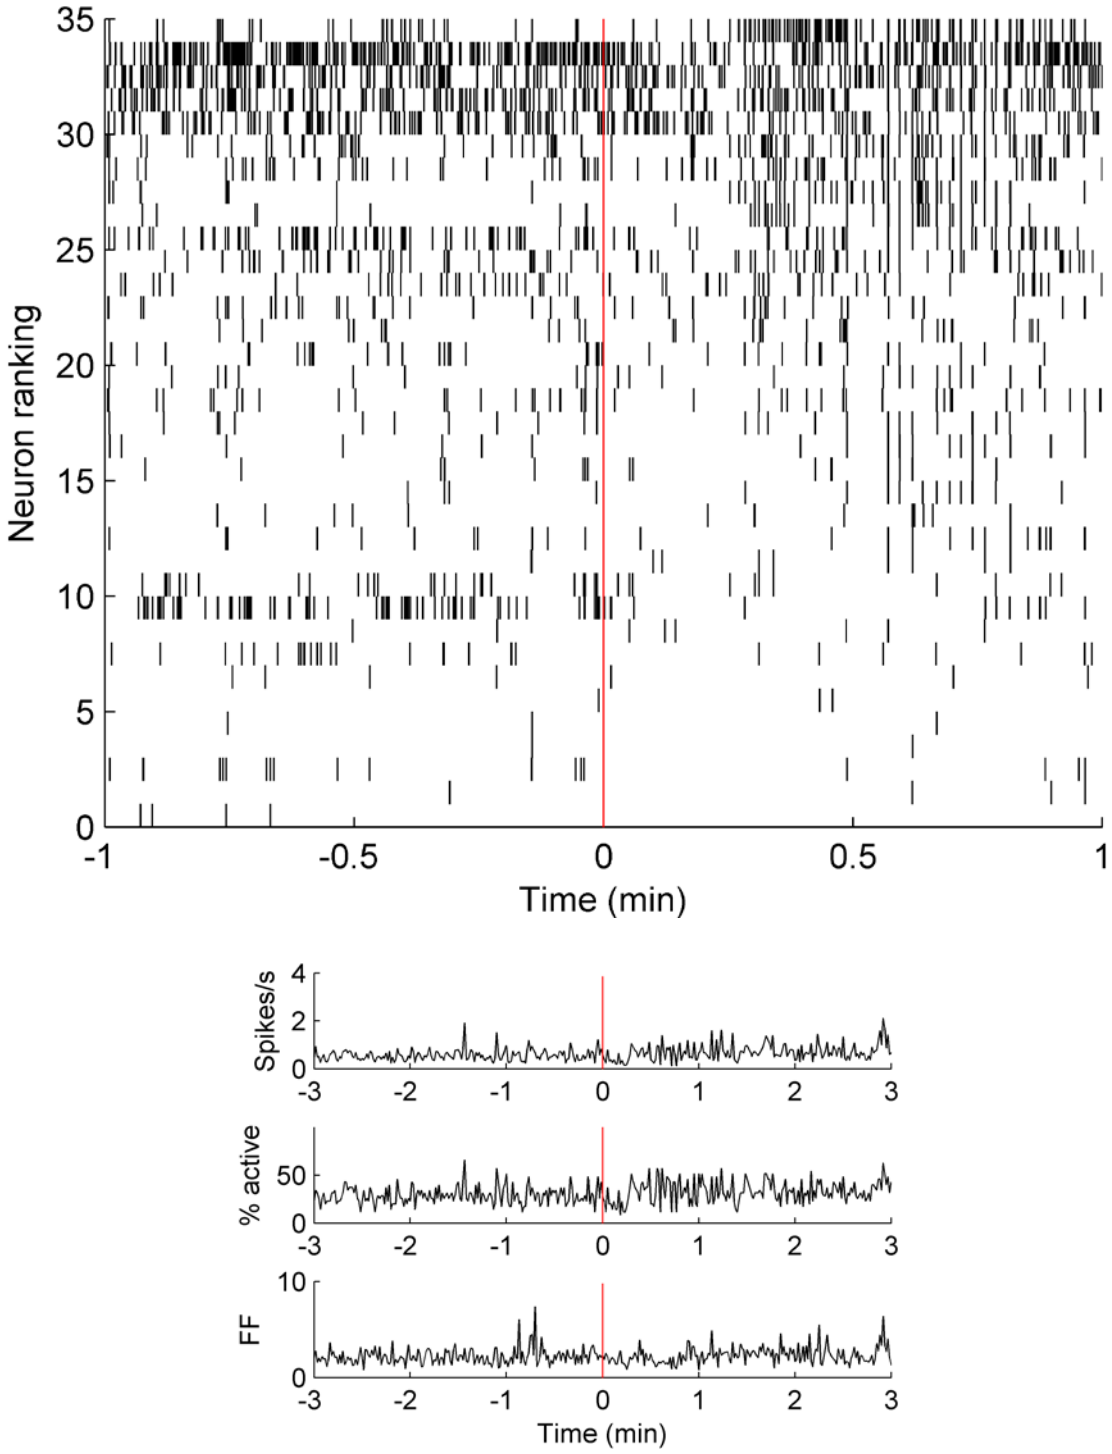

Supplementary Figure 1 (Cont.). Seizure 1 in Patient D appeared as a very mild event at the microelectrode site. It can be hardly detected on the population spike count rate, in the percentage of coactive neurons or in the Fano factor for the spike counts across the population. Nevertheless, visual inspection of the spike rasters reveals two main neuronal groups: one neuronal group with a buildup in activity (starting around 20 seconds into the seizure) and the other with a decrease during the initial 30-40 seconds. Based on ECoG analyses, the seizure lasted for 43 seconds.

Supplementary Figure 2. Examples of transient spiking activity suppression during seizure and after seizure termination. The high-pass (250 Hz – 7.5 kHz) filtered electric potential recorded at electrodes 32, 41, 42 and 47. Larger spikes in each plot correspond to units 32-1, 41-1, 42-1 and 47-1 (See Figure 2, main text). The dashed vertical lines show the seizure onset and termination, respectively. The horizontal white lines mark the ± 3SD confidence interval of the background noise estimated from the 'silent' period after seizure termination. Note that the units resume spiking at very high rates towards the end of the seizure. Although there is some gradual decrease in spike amplitudes prior to seizure termination, this decrease is much smaller than what would be expected from depolarization block. When the units stop spiking near the end of the seizure, spike peak-to-peak amplitudes are larger than 200 μV.

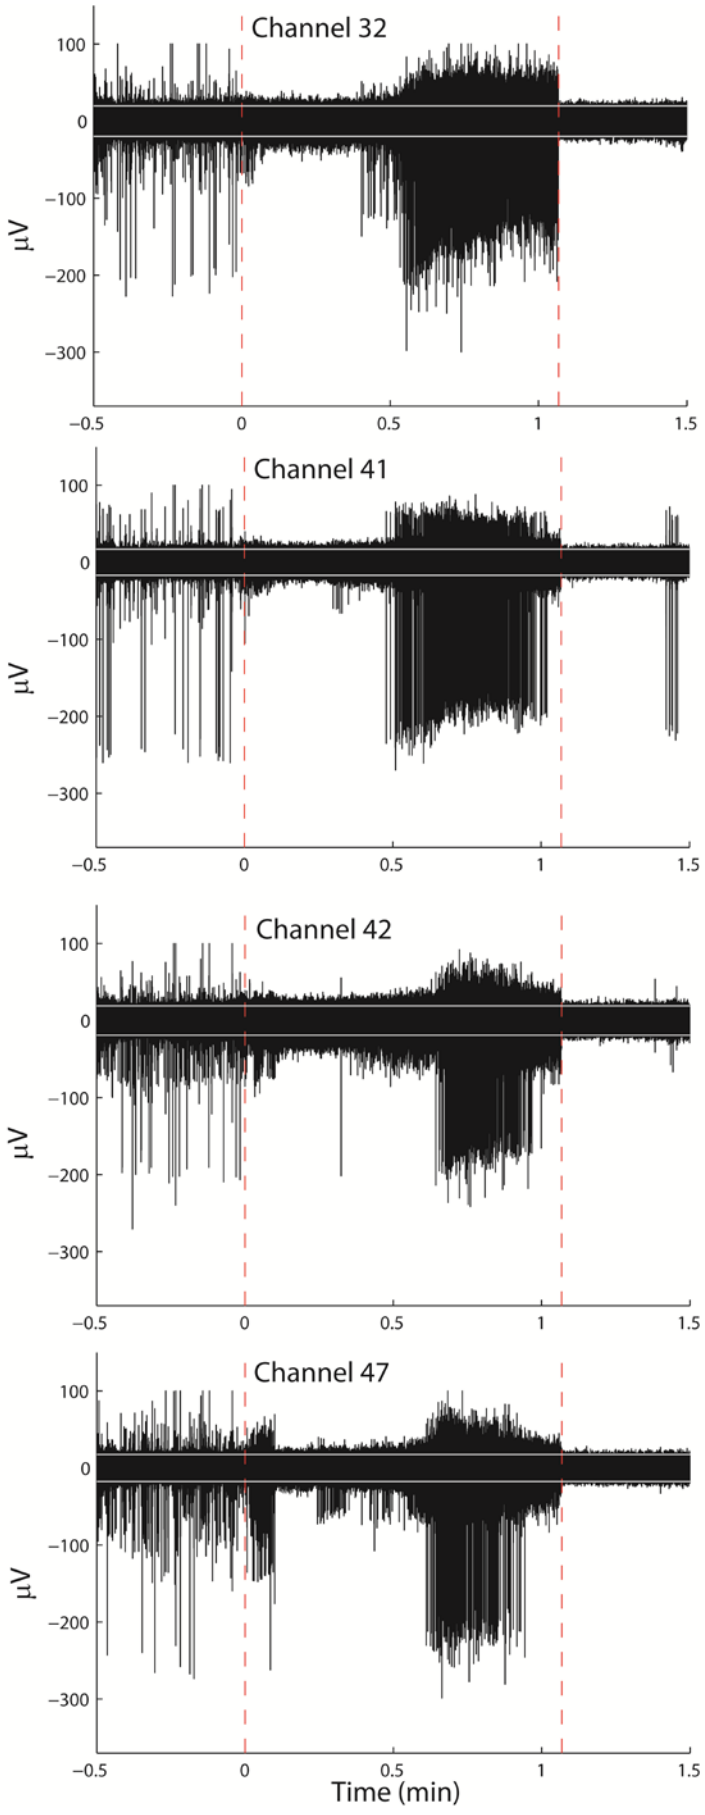

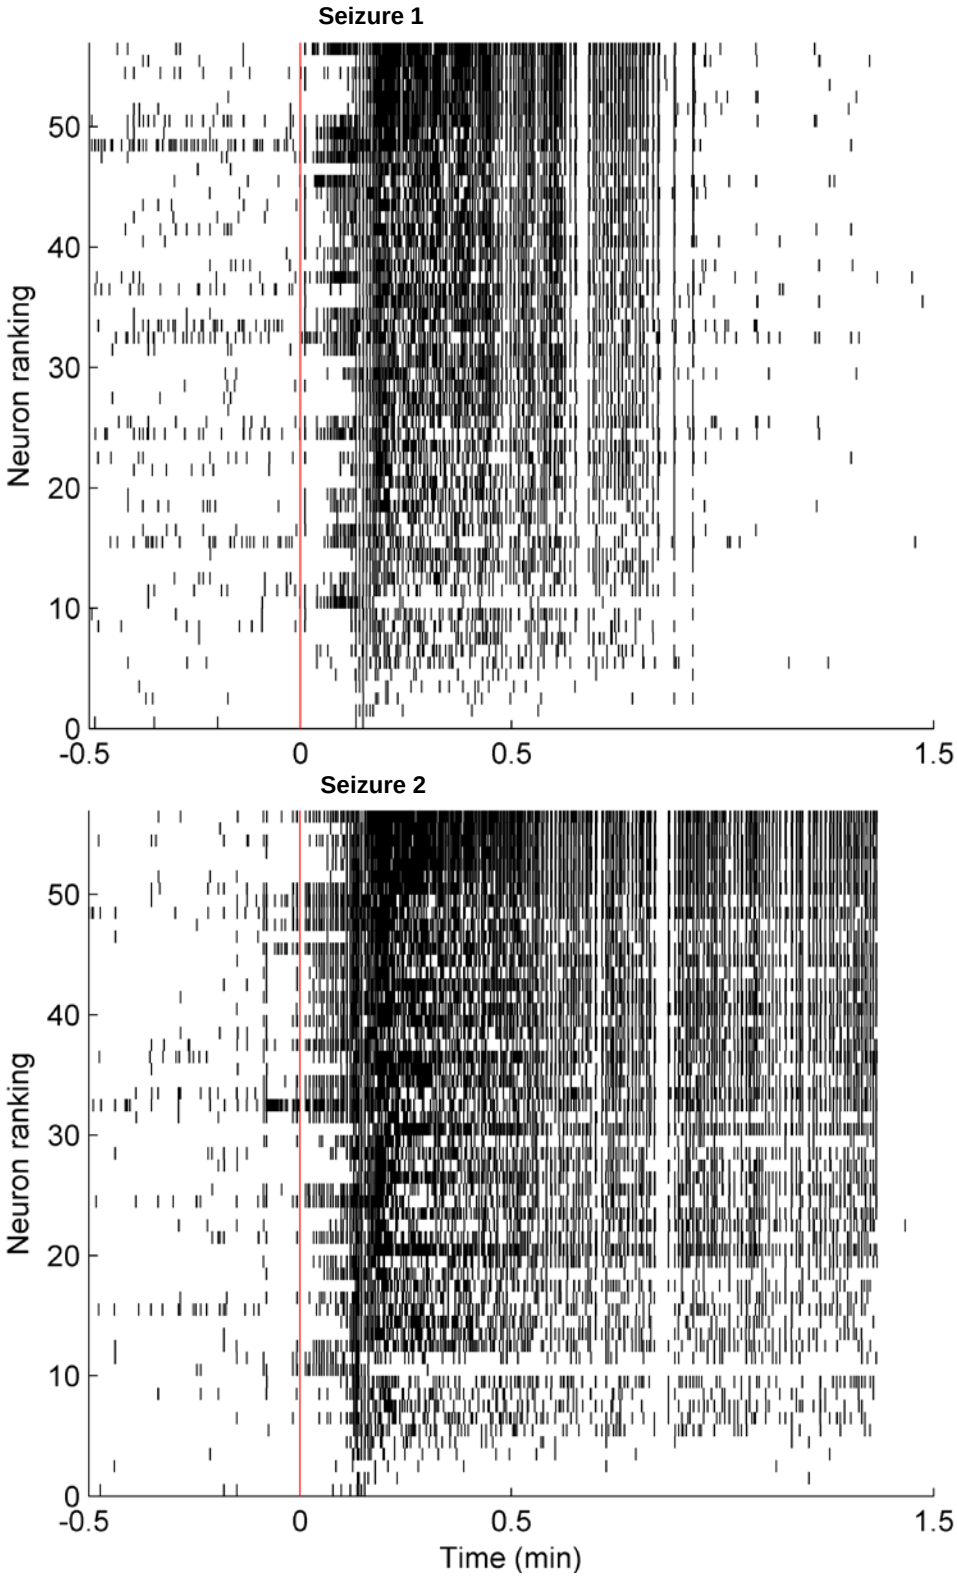

Supplementary Figure 3. Reproducibility of neuronal spiking modulation patterns across consecutive seizures (additional examples). Two seizures in Patient B occurred within ~ 1 hour allowing us to examine the reproducibility of neuronal spiking patterns across different seizures. Conventions are the same as in Figure 3, main text. Despite variations in duration, seizures 1 and 2 from Patient B show a common pattern, especially during seizure initiation. The correlation coefficient between two spike trains during the initial 30 seconds of each seizure, for each neuron and averaged across the population, was 0.72.

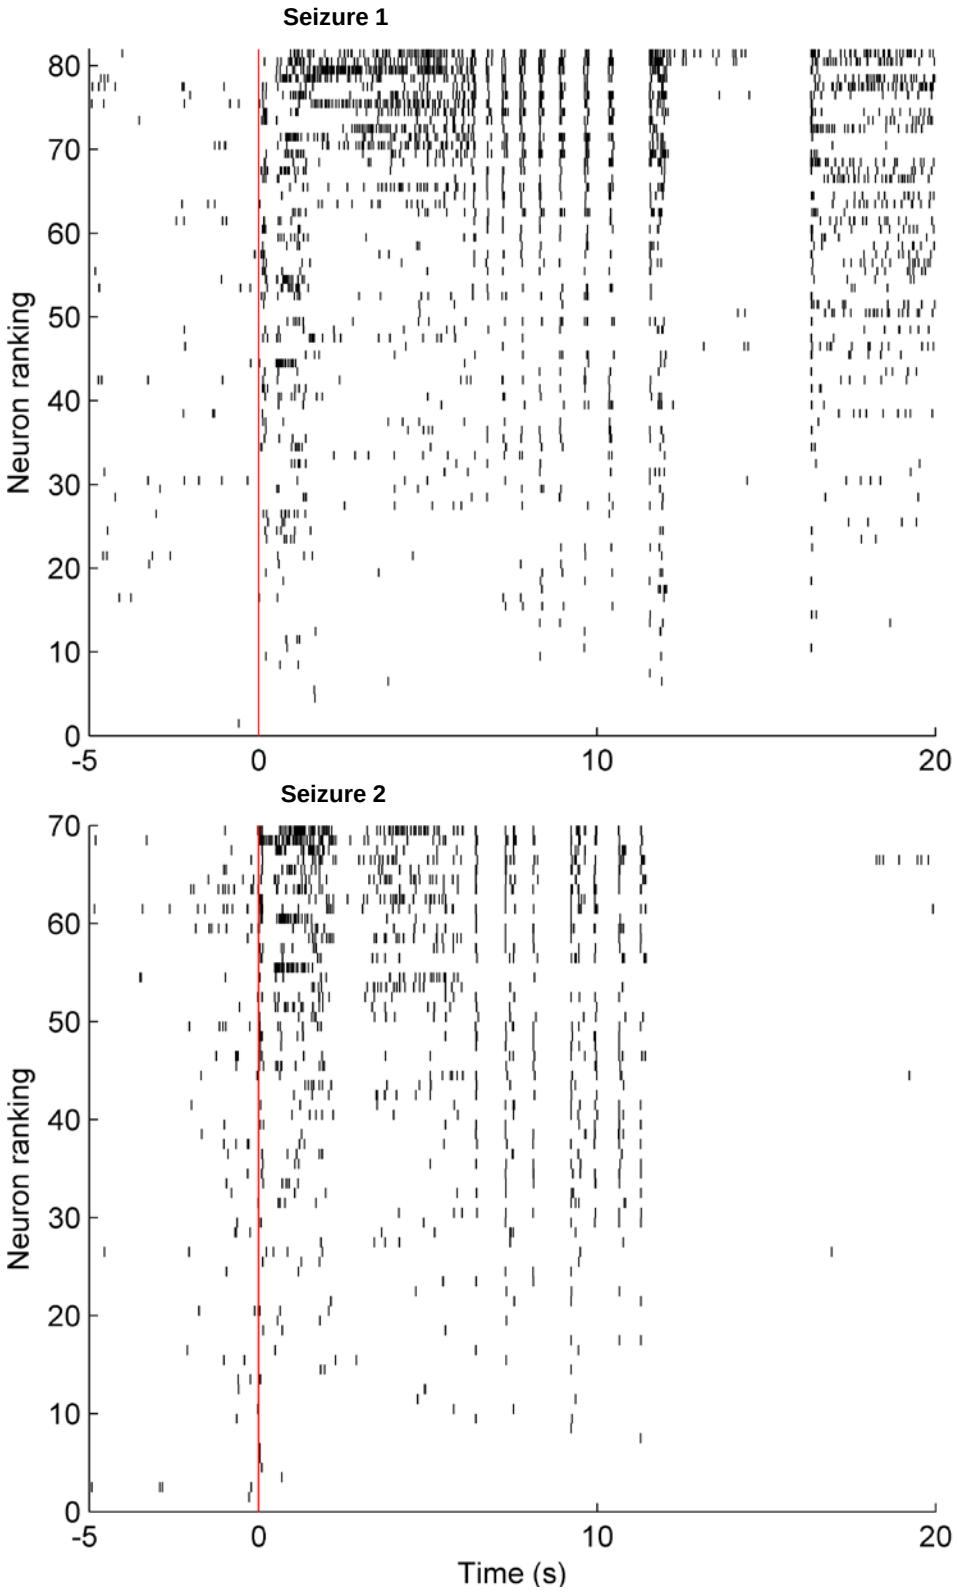

Supplementary Figure 3 (Cont.). Reproducibility of neuronal spiking modulation patterns across consecutive seizures. Seizures 1 and 2 in Patient C are shown. Even though we could not determine whether the same neurons were recorded in these two seizures because these seizures were separated by several hours, some level of reproducibility of a general modulation pattern over the entire neuronal ensemble can be observed. According to ECoG inspection, both electrographic seizures lasted ~ 11 seconds. (Note that each population raster contains a different number of neurons. In contrast with Fig. 3, main text, and the two

population rasters for Patient B in the first part of this Supp. Fig., neurons with the same ranking might relate to different single units.)

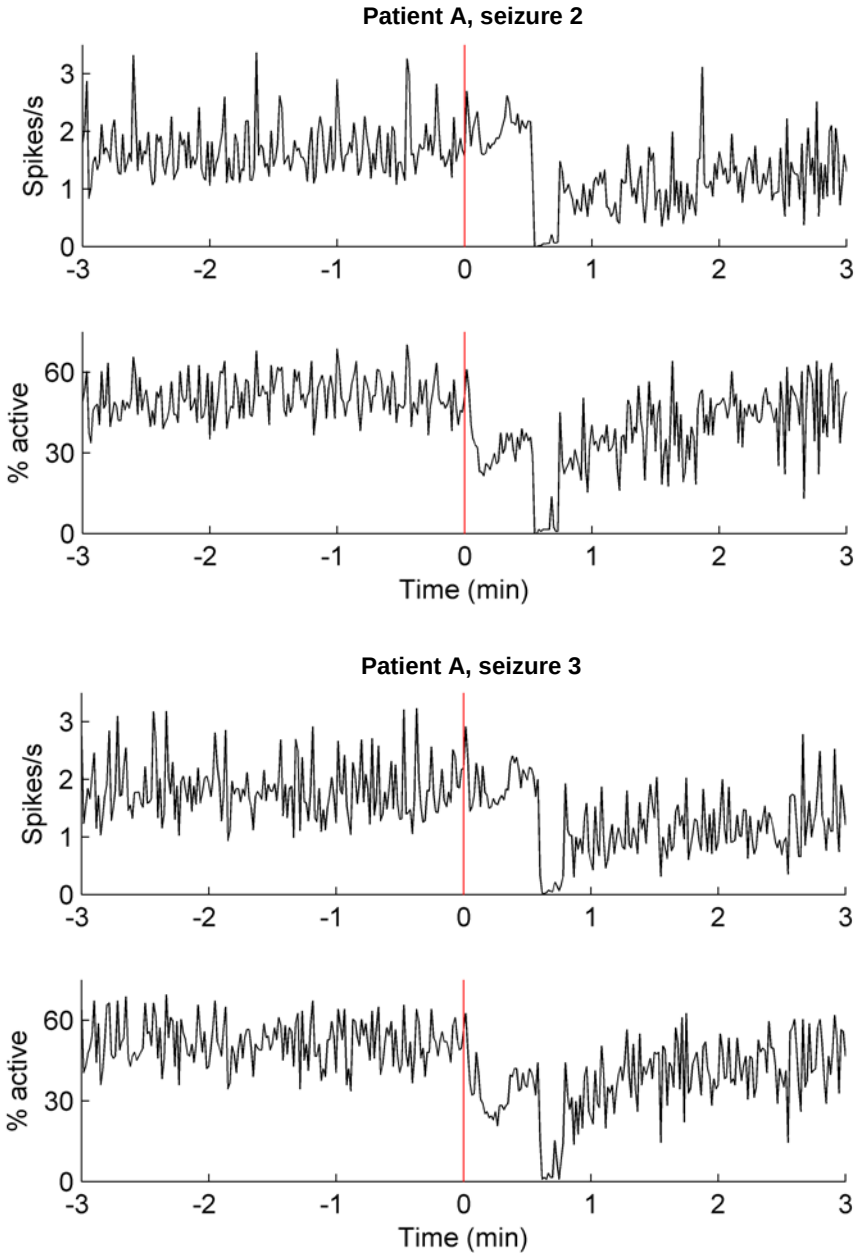

Supplementary Figure 4. Fraction of active neurons in seizures 2 and 3 (Patient A). The top plot in each panel shows the spike rate, computed from spike counts in 1-second time bins, averaged across the ensemble of neurons. Bottom plot in each panel: in contrast to seizure 1 (Fig. 1, main text), where the fraction of active neurons transiently increased towards the end of the seizure, this fraction remained below the preictal level throughout seizures 2 and 3. (Percentage of active neurons was computed on a 1-second time scale, i.e. a neuron was considered active once if spiked at least once in a given 1-second time bin.)

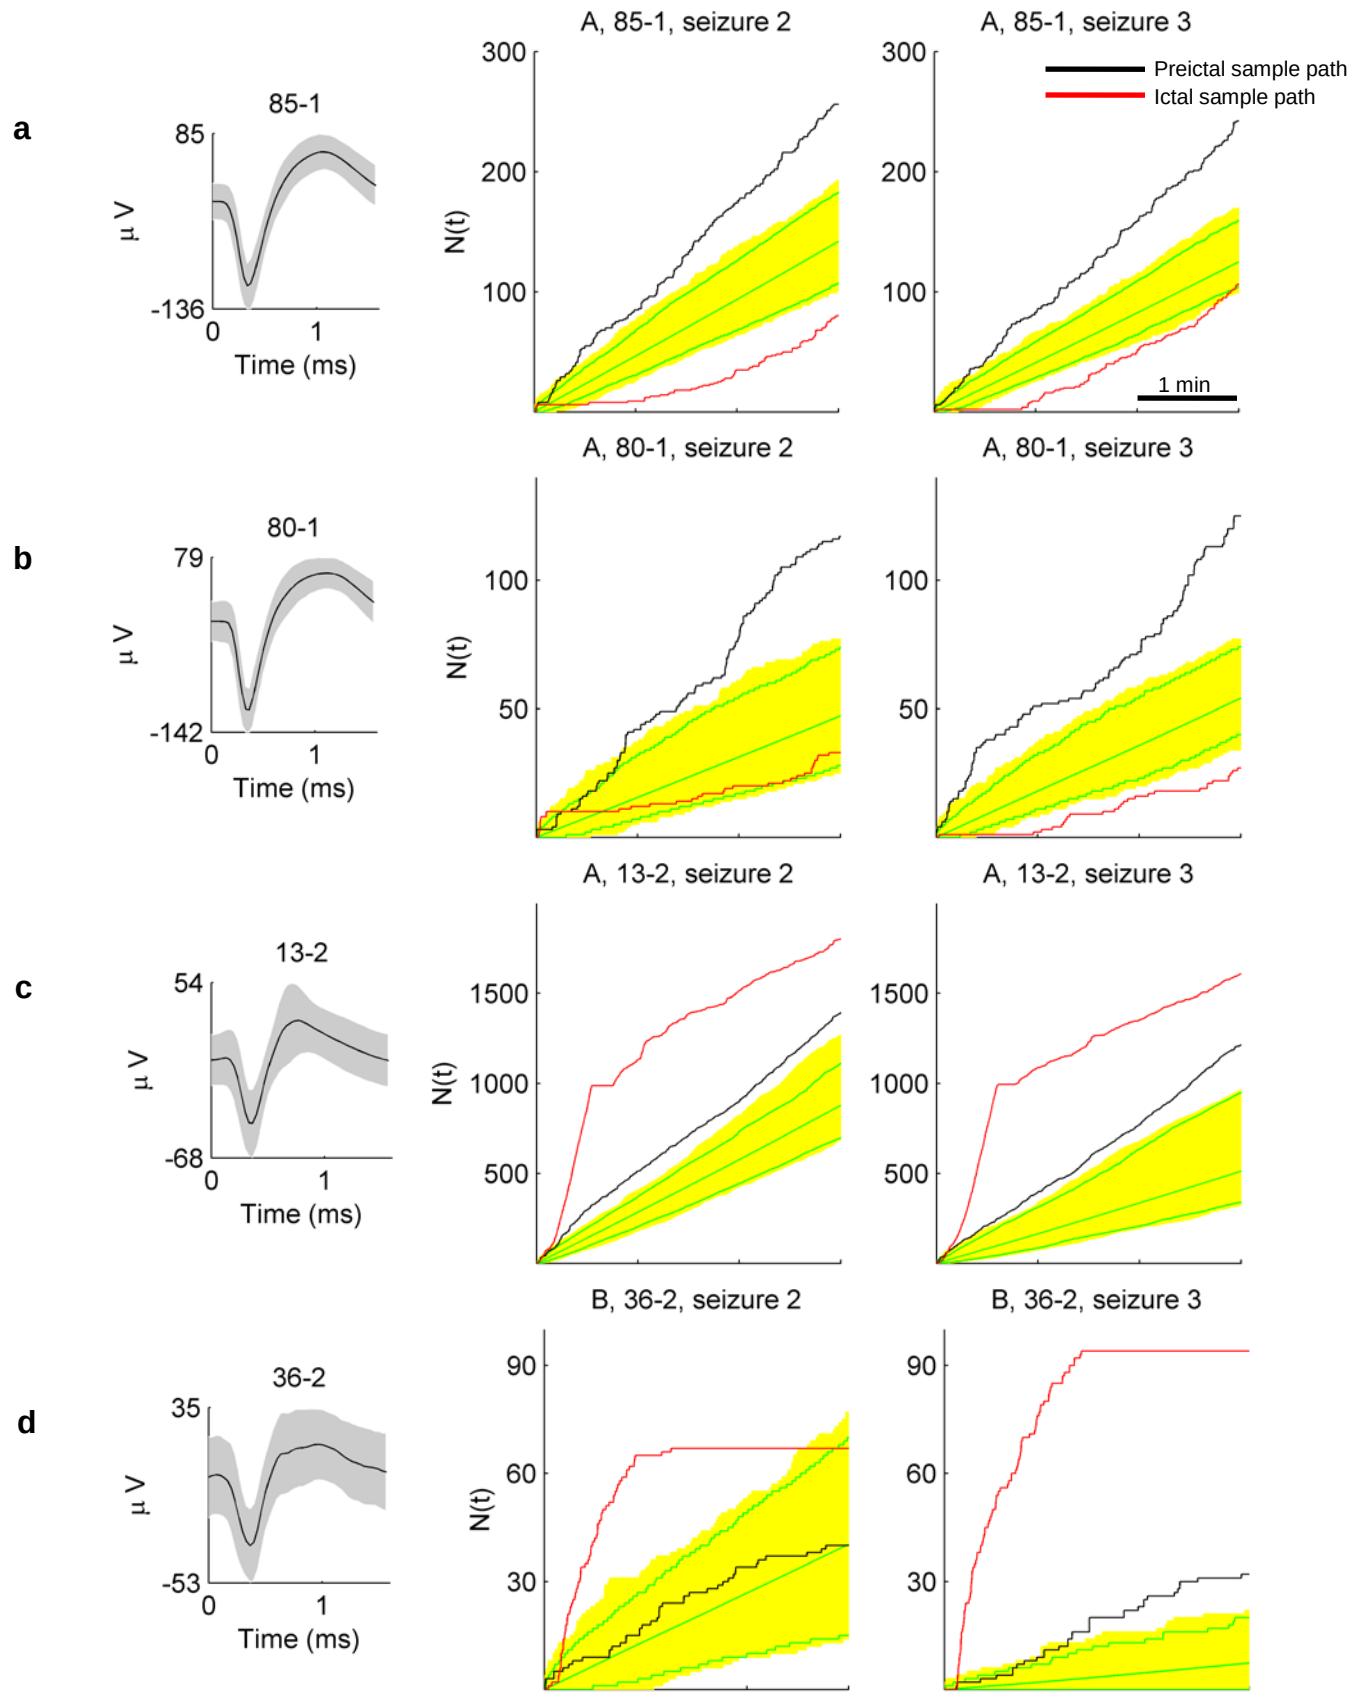

Supplementary Figure 5. Sample path deviations across different seizures. The two consecutive seizures (within ~ 1 hr.) in Patients A and B allowed us to examine the reproducibility of sample path deviations across

different seizures. (**a-c**) Despite variability in preictal and ictal activity, several neurons showed consistent preictal and ictal deviations with respect to a preceding interictal period (Patient A). (**d)** An example neuron from Patient B which showed a preictal deviation in the second seizure, but not in the first. Black and red curves in the middle and left column plots refer to preictal and ictal sample paths, respectively. The yellow region covers the range of interictal sample paths. The green curves denote the mean sample path and the 95% confidence interval.

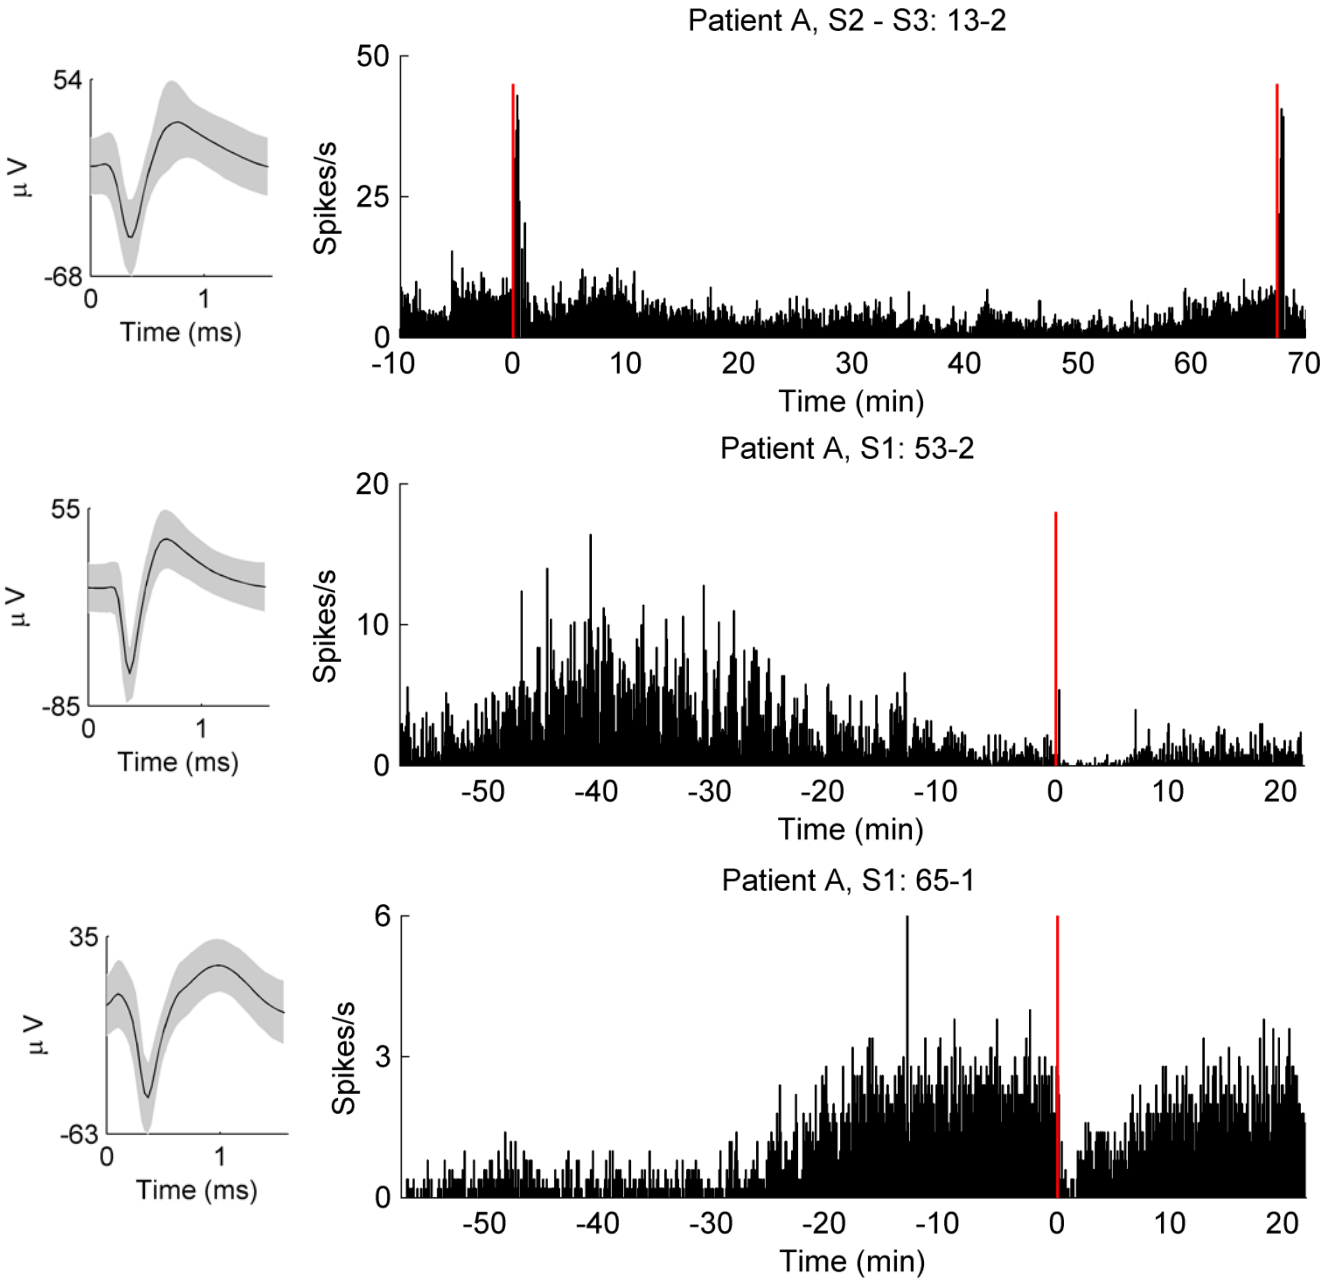

Supplementary Figure 6. Examples of slow, strong modulations in neuronal spiking rates. Spike rates computed on 5-second time bins show long term and, in one case (top plot), reproducible changes in spiking rates in neurons long in advance of the seizure onset. Seizure onsets are marked by vertical red lines.

Supplementary Figure 7. Features of neuronal spiking during interictal periods: mean firing rates, ISI coefficient of variation (CV) and bursting rate. A 30-minute long period, preceding a defined 3-minute preictal period, was chosen for characterizing basic spiking properties of the individual neurons. (A1: Patient A, seizure 1.)

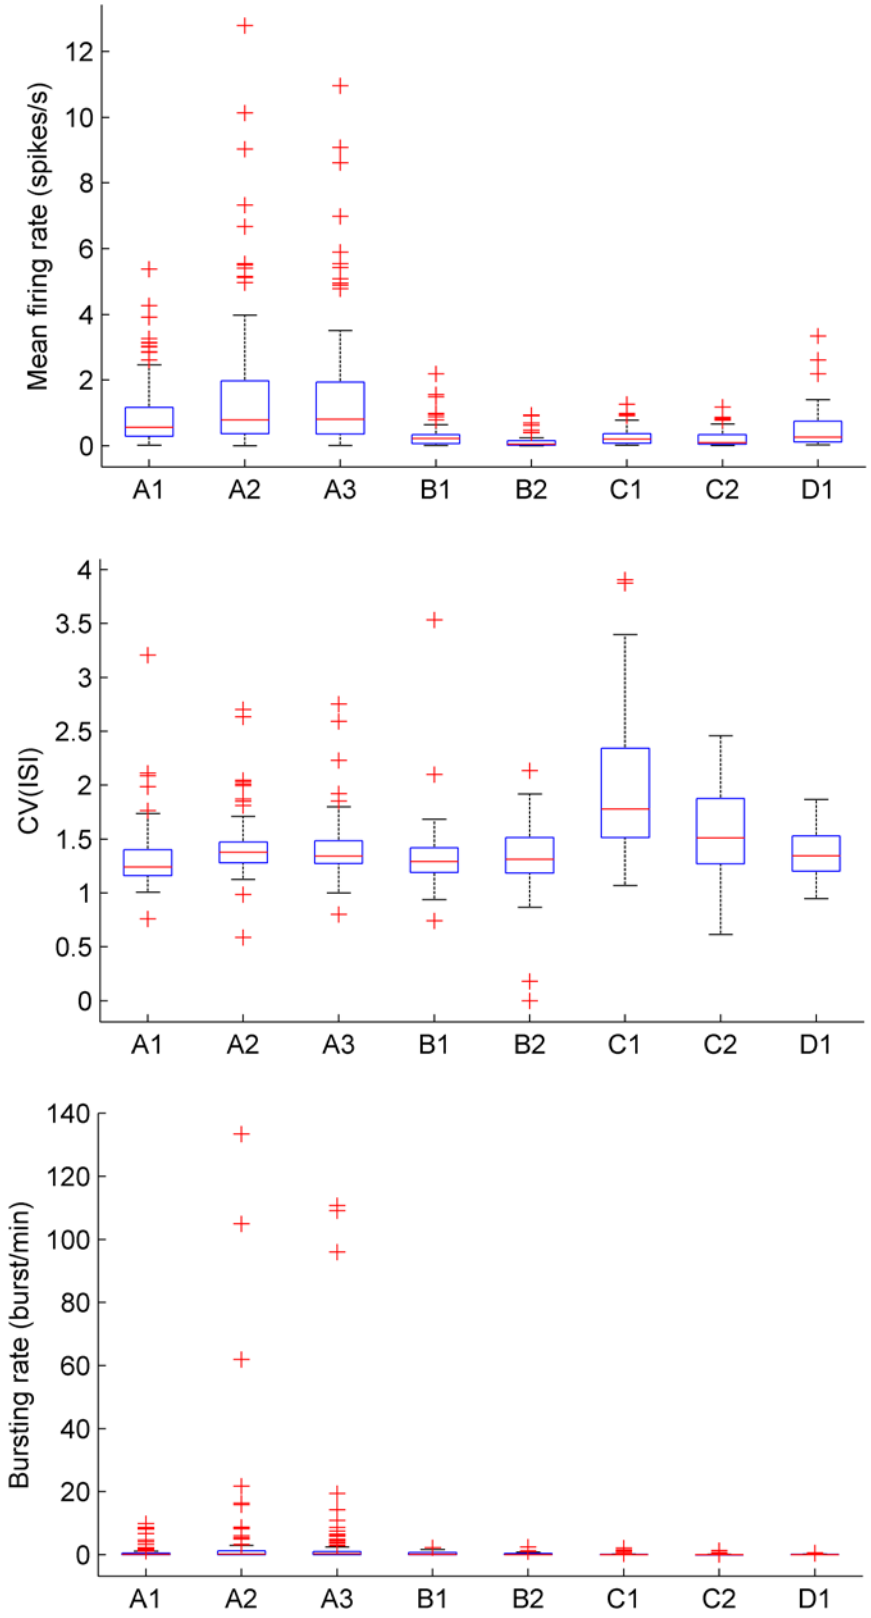

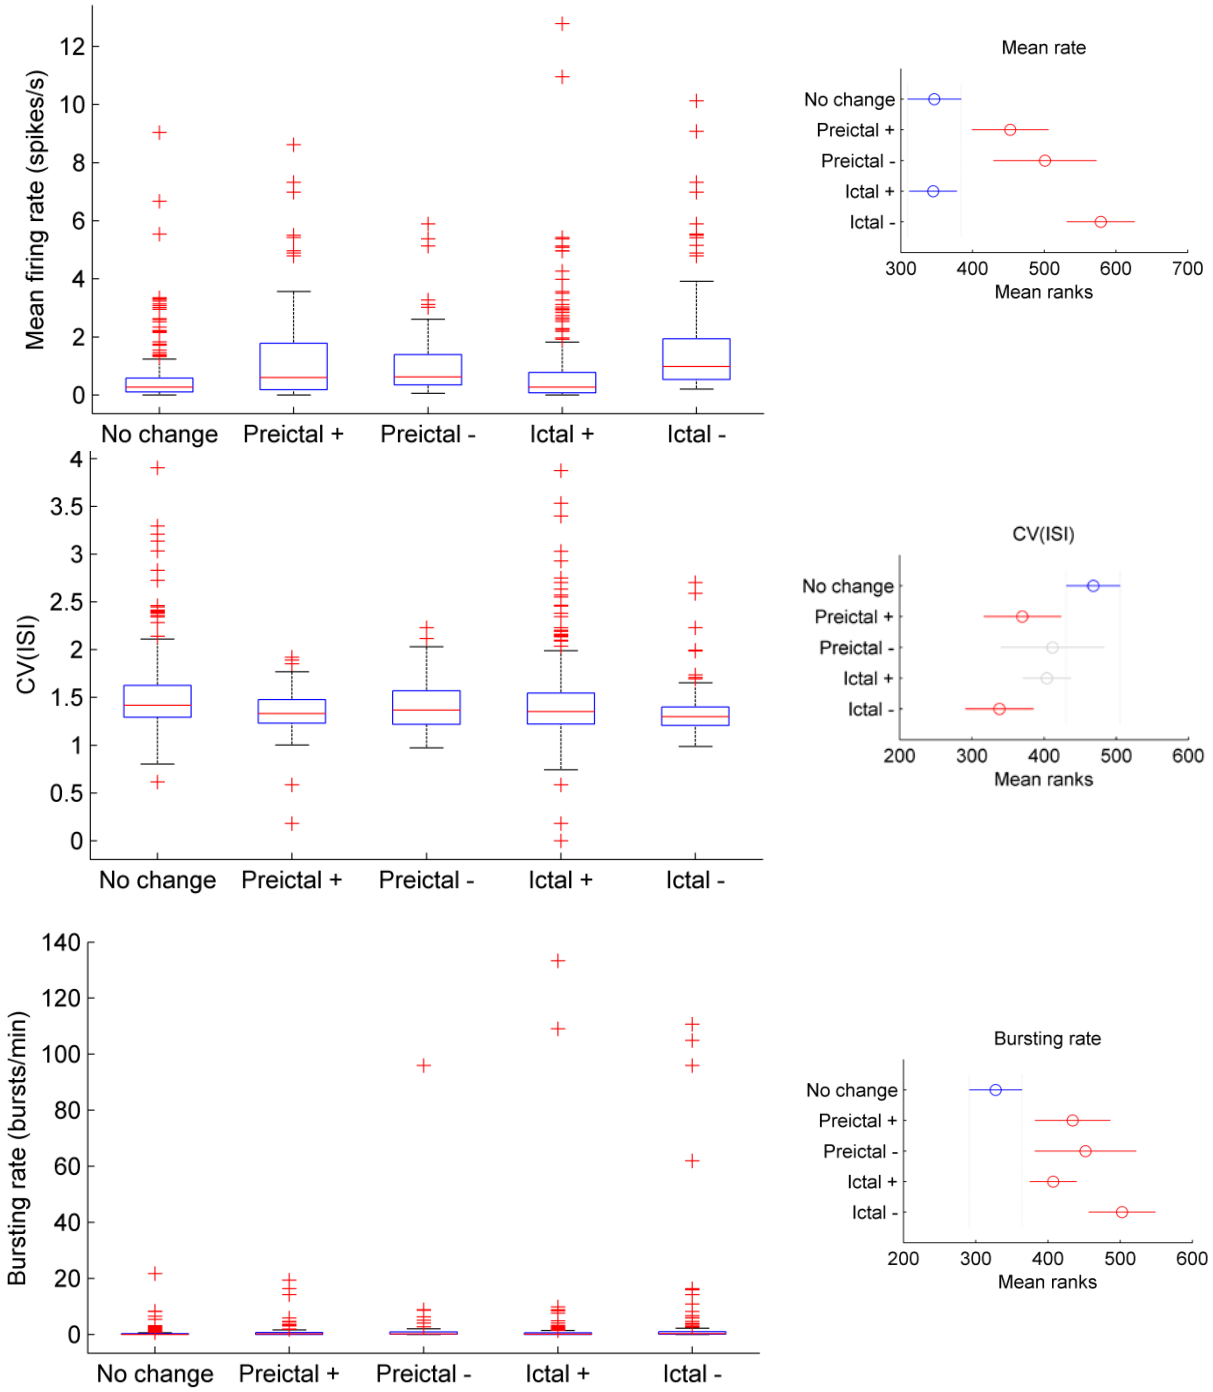

Supplementary Figure 8. Interictal features (mean firing rate, CV, bursting rate) vs. preictal and ictal modulation. Five main modulation types (including "no change") based on data from all patients and seizures were examined. Right panels show a summary from the Kruskal-Wallis test. Here, we focus in the comparison between "No change" and the other groups; in this case, "red bars" in the right panels represent significant differences ($P < 0.01$, with correction for multiple comparisons based on "honestly significant differences" (Tukey-Kramer correction)). The most noticeable difference was that neurons that showed a preictal or ictal modulation had higher bursting rates than neurons that did not show any type of change. Plus and minus signs indicate increase and decrease in firing rat, respectively.

Nature Neuroscience: doi:10.1038/nn.2782

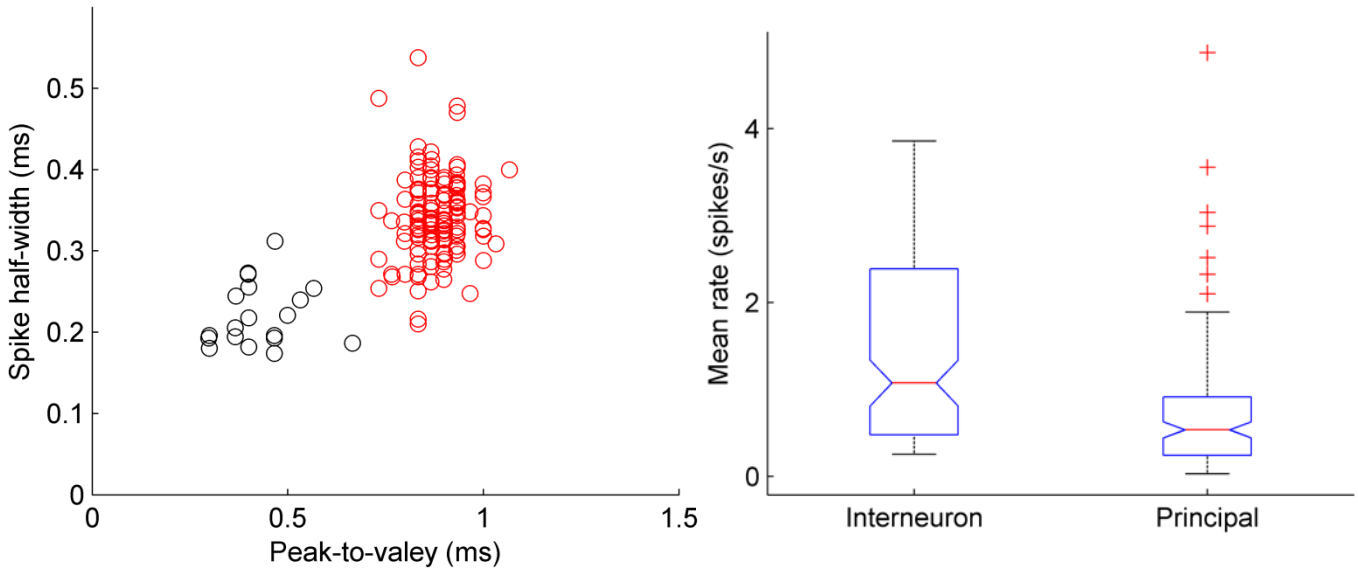

Supplementary Figure 9. Neuronal subtypes: putative principal cells and interneurons. (Patient A, interictal period preceding seizure 1; similar clustering was observed using the Kruskal-Wallis test for difference between interneuron and pyramidal cell groups (1 and 2, respectively), $P < 0.001$. Similar results were obtained for interictal periods preceding seizures 2 and 3.)

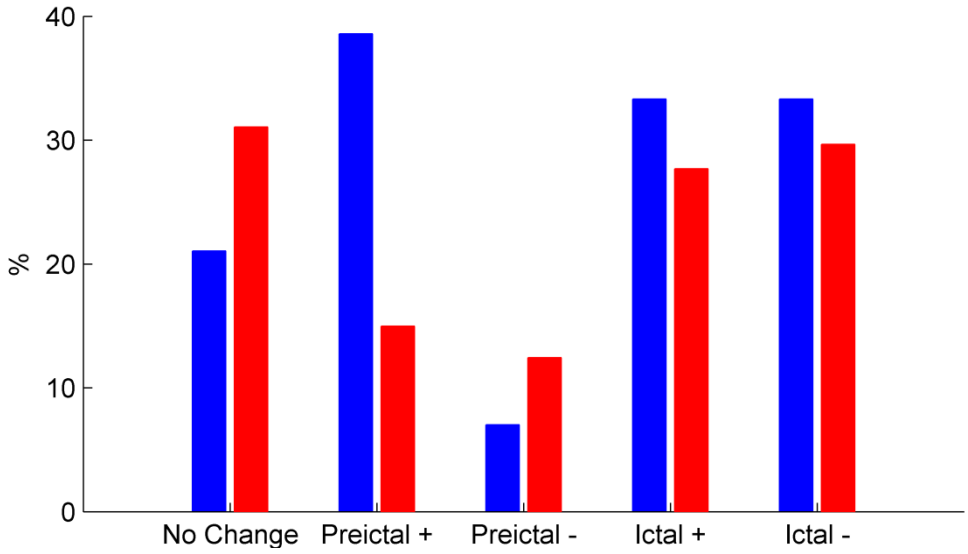

Supplementary Figure 10. Spiking rate modulation in interneurons (blue) and principal neurons (red). 21% and 32% of interneurons and principal cells, respectively, showed no modulation. The remaining neurons in both groups were distributed non-uniformly across the 4 main modulations types (preictal and ictal increase or decrease in spiking rate). In particular, fewer interneurons and principal neurons showed a preictal decrease (Chi-square test, $P < 10^{-4}$ and $P < 10^{-3}$, respectively, with Bonferroni correction for multiple comparisons). In addition, the fraction of recorded interneurons that showed a preictal increase was significantly larger than the corresponding fraction of principal neurons ($P < 10^{-6}$). (Total number of neuronal recordings: 411; 57 interneuron recordings and 354 principal cell recordings.) Plus and minus signs indicate increase and decrease in spiking rates, respectively.

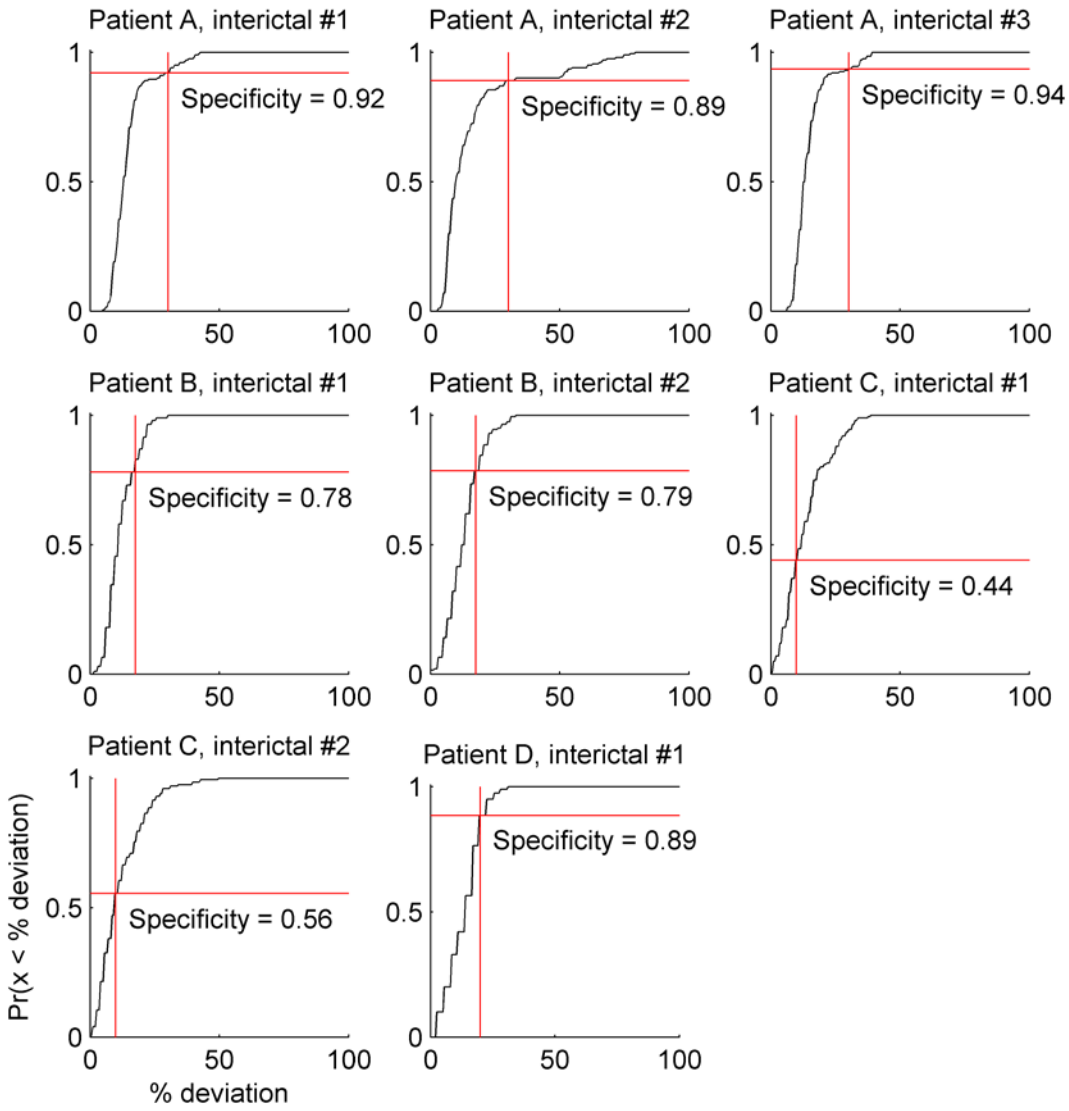

Supplementary Figure 11. Specificity of detected sample path deviations. We attempted to estimate the specificity of detected sample path deviations. Specifically, we estimated the probability of observing a percent change, i.e. a percentage of deviations across the neuronal population during interictal periods, higher than a given threshold. This threshold was set to the average percentage of preictal deviations across the examined seizures for a given patient. For example, the threshold for Patient A was obtained by computing the mean of the percentage of neurons that showed a preictal deviation across the 3 seizures. (These percentages are shown for each patient and seizure in Fig. 5, main text.) Note that this simple algorithm did not take into account either the magnitude of the deviations or which specific neuron deviated, only whether a deviation had occurred. A deviation was defined as before: every time a sample path fell outside the range of a given distribution, a deviation was detected. We randomly selected a 3-minute segment from a given interictal period and computed the corresponding (target) sample path for each neuron. A non-overlapping 30-minute period was then selected and 3-minute paths were sampled following the same procedure as in Figure 4 (main text). Based on these sample paths, a sample path distribution was computed for each neuron separately. We then asked whether the target sample path from a given neuron deviated from the corresponding distribution, i.e. if at any time the sample path fell outside the range of the distribution. After doing the same for all of the neurons, the percentage of detected deviations across the entire population and for that particular random sample was computed. We repeated the just described procedure two hundred times, and then derived a distribution of the percentages of deviations across the recorded neuronal population. Each plot above shows the corresponding cumulative distribution function (cdf) for each patient and a specific interictal period. For some of the patients, multiple

different interictal periods were examined. The vertical line marks the percentage threshold. The cdf at this threshold provided a preliminary estimate of the specificity of a seizure detection algorithm based on whether the percentage of deviation across the population was equal or above the defined threshold. For example, in the top left plot (Patient A, interictal #1), 92% of the random samples showed a percentage of deviations smaller than the mean percentage observed during the true preictal periods.

Supplementary Table I. Percentage of neuronal recordings showing preictal and ictal sample path deviations from interictal activity: preictal (rows) vs. ictal (columns) modulation (sample path deviation analysis).
0: no change; ↑: increase; ↓: decrease; ↑↓: transient increase, followed by a transient decrease.

All 4 patients, 712 neuronal recordings

	0	↑	↓	↑↓	↓↑
0	35.4	37.0	7.3	0.0	1.0
↑	3.6	6.6	1.3	0.1	0.2
↓	4.3	1.8	1.3	0.0	0.1

Patient A (3 seizures: 149, 131 and 131 neuronal recordings)

	0	↑	↓	↑↓	↓↑
0	29.7	18.7	17.5	0.0	4.1
↑	6.8	5.1	5.1	0.5	0.7
↓	4.9	4.2	2.2	0.0	0.5

Patient B (2 seizures, 57 and 57 neuronal recordings)

	0	↑	↓	↑↓	↓↑
0	2.6	79.8			
↑		15.8			
↓		1.8			

Patient C (2 seizures, 57 and 57 neuronal recordings)

	0	↑	↓	↑↓	↓↑
0	46.1	44.1			
↑	2.0	2.6			
↓	3.9	1.3			

Patient D (1 seizure, 35 neuronal recordings)

	0	↑	↓	↑↓	↓↑
0	62.8	5.7	11.4		
↑	5.7	2.9			
↓	8.6		2.9		