Published in final edited form as:

Neural Comput. 2012 October ; 24(10): 2655–2677. doi:10.1162/NECO_a_00343.

Predicting Single-Neuron Activity in Locally Connected Networks

Feraz Azhar and

Department of Neurosurgery, Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02115, U.S.A., and Department of Neurosurgery, The Johns Hopkins University School of Medicine, Baltimore, MD 21287, U.S.A

William S. Anderson

Department of Neurosurgery, The Johns Hopkins Hospital, Baltimore, MD 21287, U.S.A

Feraz Azhar: fazhar1@jhmi.edu; William S. Anderson: wanders5@jhmi.edu

Abstract

The characterization of coordinated activity in neuronal populations has received renewed interest in the light of advancing experimental techniques that allow recordings from multiple units simultaneously. Across both in vitro and in vivo preparations, nearby neurons show coordinated responses when spontaneously active and when subject to external stimuli. Recent work (Truccolo, Hochberg, & Donoghue, 2010) has connected these coordinated responses to behavior, showing that small ensembles of neurons in arm-related areas of sensorimotor cortex can reliably predict single-neuron spikes in behaving monkeys and humans. We investigate this phenomenon using an analogous point process model, showing that in the case of a computational model of cortex responding to random background inputs, one is similarly able to predict the future state of a single neuron by considering its own spiking history, together with the spiking histories of randomly sampled ensembles of nearby neurons. This model exhibits realistic cortical architecture and displays bursting episodes in the two distinct connectivity schemes studied. We conjecture that the baseline predictability we find in these instances is characteristic of locally connected networks more broadly considered.

1 Introduction

Quantifying the responses of collections of neurons to the variety of stimuli they are exposed to lies at the heart of deciphering the neural code. As one example, the role of coordinated activity in networks of cortical neurons in particular, represents a natural starting point for the exploration of how neurons encode information relevant to behavior and cognition. Of late, progress has been made toward understanding the responses of small collections of neurons involved in both lower-level sensory processing and in primary and associative cortical areas as a result of advances in experimental techniques that make simultaneous recordings from nearby populations of neurons possible (Segev, Goodhouse, Puchalla, & Berry II, 2004; Gunning et al., 2005; Suner, Fellows, Vargas-Irwin, Nakata, &

Donoghue, 2005). These experimental successes have been coupled with theoretical progress toward understanding collective behavior through the analysis of minimal statistical mechanics-based models (Schneidman, Berry II, Segev, & Bialek, 2006; Shlens et al., 2006; Tang et al., 2008), as well as point process models defined explicitly through maximum likelihood techniques (Brillinger, 1988; Chornoboy, Schramm, & Karr, 1988; Martignon et al., 2000; Okatan, Wilson, & Brown, 2005; Nykamp, 2007; Pillow et al., 2008; Truccolo, Hochberg, & Donoghue, 2010).

In the case of models based on statistical mechanics considerations, joint probability densities of networks of neurons can be constructed assuming maximum entropy distributions consistent with no higher than pairwise correlations (for example), as neurons respond to naturalistic external stimuli (Schneidman et al., 2006) and nonnaturalistic external stimuli (Schneidman et al., 2006; Shlens et al., 2006), as well as for neurons undergoing spontaneous activity (Schneidman et al., 2006; Tang et al., 2008). The resulting models suggest that pairwise interactions capture significant fractions of the information contained in the states adopted by the network. In the case of the point process models, including the intrinsic and extrinsic spiking histories for up to 100 ms, for roughly 20 to 200 neurons, allows one to successfully predict the spiking activities of single neurons in the experimentally measured ensemble (Truccolo et al., 2010). Notably, in this case, predictability is determined when the external stimuli involve performance of behavioral tasks.

Coupled together, these results suggest the existence of strongly coordinated activity in small networks of neurons across a broad range of neuronal preparations in a variety of stimulus paradigms. We investigate whether this coordinated activity can be accounted for by random fluctuations in locally connected model networks. We make progress toward addressing this question using point process modeling in the style of Truccolo et al. (2010). We show that in a recently developed computational model of cortex (Anderson, Kudela, Weinberg, Bergey, & Franaszczuk, 2009; Anderson, Kudela, Cho, Bergey, & Franaszczuk, 2007), there exists significant predictive power in small, randomly sampled ensembles of neurons driven solely by random background inputs. These inputs drive coordinated activity, which we measure through predictability of single-neuron spiking, and we conjecture that this baseline predictability, which in effect constitutes a lower bound to predictability, is a feature of locally connected networks considered quite generally.

2 The Computational Model

We institute our prediction scheme on a recently developed spiking neural network simulation with realistic cortical architecture (Anderson et al., 2009, 2007), which displays network bursting dynamics in response to random background inputs. The details of this model have been reported in Anderson et al. (2009, 2007). For completeness, we collect salient details relevant to our prediction scheme here, as well as in the appendix.

2.1 The Network

The version of the model we investigate here has 65,536 single-compartment neurons operating on a modified Hodgkin-Huxley scheme (Av-Ron, Parnas, & Segel, 1993),

including spherically symmetric calcium diffusion and synaptic currents (Kudela, Franaszczuk, & Bergey, 2003). The network represents a surface area of $1.6 \text{ mm} \times 1.6 \text{ mm}$ of simulated cortex, with a thickness of 2.34 mm . Neurons are arranged in a repeating lattice of 64×64 minicolumns, with center-to-center spacing between these minicolumns set at $25 \mu\text{m}$. Each minicolumn contains 16 neurons spanning seven different classes organized in cortical layers. Layer II/III contains four excitatory pyramidal cells, one inhibitory double bouquet cell, and one inhibitory basket cell. Layer IV contains one excitatory stellate cell and one inhibitory chandelier cell. Layers V and VI each contain four excitatory pyramidal cells. Intracolumnar wiring is adapted from visual cortical models and generic somatosensory models (Douglas & Martin, 2004; Nieuwenhuys, 1994). Extracolumnar wiring is based on histological data for the numbers of contacted postsynaptic cells and is generally isotropic in space within the plane of the postsynaptic cell. The spatial distributions of connections are typically flat out to 200 to $300 \mu\text{m}$, with an exponential tail that can extend to 1 mm in the case of inhibitory basket cells.

For the purpose of our prediction scheme, we chose random subsets of 20 and 40 layer II/III pyramidal cells from the central two-thirds of the model, representing single-neuron recordings from an area of about 1.1 mm^2 . Figure 1 schematically displays the relative locations of each cell in both of these ensembles.

2.2 Background Input and Network Activity

The network is driven by a spatially distributed Poisson background that affects 1% of the modeled neurons. A synaptic current injection causes these neurons to spike with high probability, with the sequence of input pulses for any one neuron constituting a Poisson process. The baseline synaptic input rate for each affected cell is 2 Hz. The background input applied to different neurons is uncorrelated (Anderson et al., 2009).

The resulting activity alternates in a random fashion between periods of bursting and quiescence. Bursting is characterized by a flat frequency spectrum. Varying the pattern of connections or sequence of background inputs produces similar random activity. Alternations in the frequency of the applied background input, or numbers of connections between cells, alters this activity in a graded fashion, with higher connectivity and higher background input frequency producing constant bursting behavior.

In this letter, we probe two separate random patterns of connectivity in the network. In the first, which we refer to as connection pattern 1, each layer II/III pyramidal cell contacts at most 178 other layer II/III pyramidal cells out to a radius of $300 \mu\text{m}$ (about 10% of all layer II/III pyramidal cells encountered within that radius). The resulting network behavior during a bursting period for both of the 20- and 40-layer II/III pyramidal cell ensembles is displayed in Figures 2a and 2b, respectively. The model run for connection pattern 1 lasted 330 s and was sampled at 1 kHz. The extremal Pearson correlation coefficient, defined as that having the greatest absolute value over all time lags between $\pm 500 \text{ ms}$, is less than 0.15 in either subset.¹

¹To be clear, we compute Pearson correlation coefficients for spike counts binned in 1 ms bins for each time lag between $\pm 500 \text{ ms}$, and then report the correlation coefficient with the greatest absolute value (Welsh, 1999).

In the second pattern of connectivity (connection pattern 2), each layer II/III pyramidal cell contacts at most 112 other layer II/III pyramidal cells out to a radius of 300 μm (about 6% of all layer II/III pyramidal cells encountered). The resulting network behavior during a bursting period is displayed in Figure 3. The model run for connection pattern 2 lasted 300 s and was also sampled at 1 kHz. The extremal correlation coefficient in either subset is less than 0.25. We note that intra- and extracolumnar wiring for both patterns of connectivity for all other pairs of cells was identical.

3 Inference and Prediction Formalism

The statistical framework within which we propose to perform single-neuron spike prediction is that of point process modeling. In summary, we are interested in writing down a statistical model describing the spiking behavior of a single neuron assuming this model is parametrically related to covariates that we *choose*. These parameters are set by optimizing the likelihood of the spike train, namely, by maximizing the joint probability distribution of the spiking activity subject to regularization constraints on the parameters. This joint probability distribution can be written in terms of the *conditional intensity function* of the process whose functional form we need to explicitly specify. We outline the pertinent results from the discrete-time point process analysis that we lean on to perform prediction. (For further details, see Truccolo, Eden, Fellows, Donoghue, & Brown, 2005; Truccolo et al., 2010; Brown, Barbieri, Eden, & Frank, 2004; Truccolo, 2010; Daley & Vere-Jones, 2003.)

3.1 Discrete-Time Point Processes

Consider an experiment of length T ranging over the time interval $(0, T]$, over which we observe spiking activity for an ensemble of $N_c + 1$ cells. We partition the interval into M equal subintervals of length Δ , $\{(t_{i-1}, t_i]\}_{i=1}^M$, choosing M such that at most, one spike falls in any one bin ($M = T/\Delta$). We are interested in constructing the joint probability distribution of the observed spike train for a single neuron in our ensemble, neuron I , say, and we assume we are free to use information from the other N_c neurons in the ensemble to help us. Let N_j^I represent the number of spikes that have occurred up to and including time t_j for neuron I (i.e., in the interval $(0, t_j]$). Then the quantity

$$\delta N_j^I := N_j^I - N_{j-1}^I \quad (3.1)$$

takes the value 1 if a spike occurred in the time interval $(t_{j-1}, t_j]$ or 0 if a spike did not occur in that interval for neuron I . This counts the number of spikes in the j th bin for neuron I . Let H_j contain the history of neural spiking for all $N_c + 1$ neurons up to but not including the j th bin, and denote the conditional intensity for neuron I in the j th bin by $\lambda_I(t_j | \vec{\theta}, H_j)$, where $\vec{\theta}$ contains the parameters we define the conditional intensity to depend on. Then one can write the probability of neuron I firing a spike in the j th bin to $O(\Delta)$, as (Truccolo et al., 2005, 2010; Brown et al., 2004; Truccolo, 2010),

$$P(\delta N_j^I = 1 | \vec{\theta}, H_j) \approx \lambda_I(t_j | \vec{\theta}, H_j) \Delta. \quad (3.2)$$

This allows one to write down the joint probability distribution for the spike train for neuron I (i.e., the likelihood) over the observed time interval $(0, T] \equiv (t_0, t_M]$, as

$$P_I(\vec{\theta}) = \exp \left(\sum_{j=1}^M \left\{ \log \left[\frac{\lambda_I(t_j | \vec{\theta}, H_j) \Delta}{1 - \lambda_I(t_j | \vec{\theta}, H_j) \Delta} \right] \delta N_j + \log[1 - \lambda_I(t_j | \vec{\theta}, H_j) \Delta] \right\} \right). \quad (3.3)$$

3.2 Logistic Conditional Intensity Function

What remains to be specified is the precise functional form of the conditional intensity function $\lambda_I(t_j | \vec{\theta}, H_j)$, together with the way in which any regularization parameters we may want to use enter into the formalism.

We follow the spiking histories of the $N_c + 1$ neurons in $\Delta = 1$ ms bins and consider the influence of previous spikes over a timescale of 100 ms (as in Truccolo et al., 2010).

Denoting the spike train of cell I up to but not including the j th bin by $\vec{\sigma}_j^I$, the history of our ensemble of $N_c + 1$ neurons is $H_j = \{\vec{\sigma}_j^I, \{\vec{\sigma}_j^\nu\}_{\nu=1}^{N_c}\}$. We project the intrinsic spike train (cell I) onto $N = 10$ raised cosine functions $\{\vec{B}^m\}_{m=1}^N$ of 100 ms in length and each of the external cells' spike trains onto $N^+ = 4$ such functions $\{\vec{C}^n\}_{n=1}^{N^+}$ (Pillow, Paninski, Uzzell, Simoncelli, Chichilnisky, 2005; Pillow et al., 2008; Truccolo et al., 2010). We introduce a parameter for each of the raised cosine functions and propose to utilize a logistic form for the conditional intensity (Zhao & Iyengar, 2010). Putting this together gives

$$\log \left(\frac{\lambda_I(t_j | \vec{\theta}, H_j) \Delta}{1 - \lambda_I(t_j | \vec{\theta}, H_j) \Delta} \right) = \mu + \sum_{m=1}^N \alpha_m \vec{B}^m \cdot \vec{\sigma}_j^I + \sum_{\nu=1}^{N_c} \sum_{n=1}^{N^+} \beta_{n;\nu} \vec{C}^n \cdot \vec{\sigma}_j^\nu, \quad (3.4)$$

where the free parameters in the problem $\vec{\theta} = \{\mu, \{\alpha_m\}_{m=1}^N, \{\beta_{n;\nu}\}_{n=1:N^+; \nu=1:N_c}\}$, $\{\beta_{n;\nu}\}_{n=1:N^+; \nu=1:N_c}$ number $1 + N + N_c \cdot N^+$, with μ representing a term related to the back-ground firing rate of cell I .

3.3 Raised Cosine Functions

The functions $\{\vec{B}^m\}_{m=1}^N$ and $\{\vec{C}^n\}_{n=1}^{N^+}$ introduced above partially overlap over the 100 ms of history, with one function in each of the two subsets of functions tapering off to zero just prior to 100 ms in the past. They were derived from code associated with Pillow et al. (2008); available online at http://pillowlab.cps.utexas.edu/code_GLM.html. We provide in this section one way to piece together our implementation of the code accompanying Pillow et al. (2008), as well as the equations presented there for generating these functions.

In the intrinsic history case, we construct $N = 10$ functions. Setting $\epsilon = 0.001$ and discretizing the time axis so that $t = 0, \epsilon, 2\epsilon, \dots, 0.1 - \epsilon$, gives 100 points that are used to encode the previous 100 ms of the neuron's own spiking history (in 1 ms bins). A refractory period of 2 ms is enforced using the function $B^1(t)$ (Truccolo et al., 2010), where $B^1(t) = 1$

for $t = 0, \varepsilon$, and $B^1(t) = 0$ otherwise. The remaining $N - 1$ functions are defined for $j = 2, 3, \dots, N$ as $B^j(t) = \frac{1}{2}[1 + \cos(a \log(t+c) - \phi_j)]$ for $|a \log(t+c) - \phi_j| \leq \pi$, with $B^j(t) = 0$ otherwise. Here, c is a parameter controlling the nonlinear stretching of the time axis,

$a = \pi(N-2)\{2\log(\frac{x+c}{y+c})\}^{-1}$, $\phi_j = a \log(y+c) + \frac{\pi}{2}(j-2)$, where x and y control the positions of the peaks of the last and first raised cosine bumps respectively. Two further amendments are made such that $B^j(t) = 0$ for $t = 0, \varepsilon$, again for $j = 2, 3, \dots, N$, and $B^2(t) = 1$ for $t = 2\varepsilon, \dots, y$.

The N resulting 100-element vectors $\{\vec{B}^m\}_{m=1}^N$ then correspond to the functions so defined on our discretized lattice, with the choices $c = 0.5$, $y = 0.01$, and $x = 0.08$ (set by hand). Figure 4a displays intrinsic history functions as generated through the code accompanying Pillow et al. (2008; and as directly used in the computations presented in our letter).

For the extrinsic history case, we construct $N^+ = 4$ such functions. Again setting $\varepsilon = 0.001$ and discretizing the time axis so that $t = 0, \varepsilon, 2\varepsilon, \dots, 0.1 - \varepsilon$, we form N^+ functions C^n , for $n = 1, 2, \dots, N^+$, where $C^n(t) = \frac{1}{2}[1 + \cos(a' \log(t+c') - \phi'_n)]$ for $|a' \log(t+c') - \phi'_n| \leq \pi$, with $C^n(t) = 0$ otherwise. Here, $a' = \pi(N^+ - 1)\{2\log(\frac{x'+c'}{y'+c'})\}^{-1}$, $\phi'_n = a' \log(y'+c') + \frac{\pi}{2}(n-1)$, for $n = 1, 2, \dots, N^+$. We further set $C^1(t) = 1$ for $t = y'$. The N^+ resulting 100-element vectors $\{\vec{C}^n\}_{n=1}^{N^+}$ correspond to the functions so defined, with the choices $c' = 0.5$, $y' = 0.01$, and $x' = 0.06$ (again set by hand). Figure 4b displays corresponding extrinsic history functions as generated through the code accompanying Pillow et al. (2008).

3.4 Optimization Scheme

We set the parameters $\vec{\theta}$ through a scheme effectively implementing maximum a posteriori estimation with a gaussian prior over the parameters.² In practice, one can think of this as performing optimization over an L2 regularized log likelihood, with an effective objective function $L_I(\theta, \eta, \bar{\eta})$ given by

$$L_I(\vec{\theta}, \eta, \bar{\eta}) = \log P_I(\vec{\theta}) - \eta \sum_{m=2}^N \alpha_m^2 - \bar{\eta} \sum_{\nu=1}^{N_c} \sum_{n=1}^{N^+} \beta_{n;\nu}^2, \quad (3.5)$$

where η represents a parameter controlling the regularization of the parameters related to cell I 's intrinsic history and $\bar{\eta}$ controls the regularization of the parameters related to the N_c extrinsic histories. We note that we do not regularize over parameter α_1 . This constant premultiplies $\vec{B}^1$, which is chosen to enforce the 2 ms refractory period in the model (Truccolo et al., 2010). For computational efficiency, we set the parameters η and $\bar{\eta}$ by hand, effectively probing a subspace of the support of the objective function, equation 3.5. In particular, we chose $\bar{\eta} \gg \eta$ (with $\bar{\eta} = 10^3 \eta$) to probe the regime where the external cells in the model served as perturbations to our underlying predictive capability based solely on a neuron's own intrinsic spiking history. In this way, our results represent lower bounds on

²We thank Sridevi Sarma for passing along to us a fast truncated L2 regularized binomial logistic regression solver, created by Demba Elimane Ba, MIT Department of Electrical Engineering and Computer Science, Neuroscience Statistics Research Lab (MIT Department of Brain and Cognitive Sciences), August 25, 2008.

predictive capabilities in this theoretical setting. Once we have chosen the optimal parameters in this way, one can use equation 3.2 to compute the predicted probabilities of spiking.

4 Establishing a Lower Bound on Predictability

The prediction scheme described in section 3 leads to estimates for the parameters of the problem $\vec{\theta} = \{\mu, \{\alpha_m\}_{m=1}^N, \{\beta_{n;\nu}\}_{n=1:N^+; \nu=1:N_c}\}$. Once these parameters have been set, one can compute the intrinsic and extrinsic history filters, which quantify the influence of prior spikes on the probability of spiking for any target neuron. Figure 5 displays schematic temporal filters for a single neuron out of the 20-cell subset and a single neuron out of the 40-cell subset for connection pattern 1. The parameters that determine these displayed filters are as computed for one-fold of the k -fold cross-validation schemes we used to generate results in this letter ($k = 11$ for connection pattern 1, $k = 10$ for connection pattern 2). We display the analogous plot for connection pattern 2 in Figure 6.

Applying temporal filters to our assumed form for the conditional intensity function, equation 3.4, and substituting the resulting expression into that for the probability of spiking, equation 3.2, gives predicted conditional probabilities for spiking given recordings over the previous 100 ms of activity in the ensemble. An example of the conditional probabilities thus generated for both connection patterns is presented in Figures 7 and 8.

Applying thresholds to the conditional probabilities calculated as in Figures 7 and 8, we computed cross-validated receiver operator characteristics (ROC) curves (Fawcett, 2006; Truccolo et al., 2010). ROC curves for two cells are shown in Figures 9a, 9b, 10a, and 10b, taken from both the 20- and 40-cell subsets for both connection patterns.

As in Truccolo et al. (2010), we measure predictive performance by computing the area under the ROC curve (AUC), as well as the chance level AUC (denoted by AUC*), obtained here by randomly shuffling the spike trains of the cells under consideration, and then define predictive power as $2(\text{AUC} - \text{AUC}^*)$. A predictive power of 0 signals chance-level performance, whereas a predictive power close to 1 signals almost perfect performance. Histograms for the 20-cell case for both connection patterns are shown in Figures 9c and 10c and for the 40-cell case in Figures 9d and 10d. Typical AUC* values were about 0.50. We distinguish between the case where the optimization is done solely with the intrinsic spiking history of the target cell (shown in gray in Figures 9 and 10) and the case where the optimization is done with the inclusion of the remaining cells in the ensemble (shown in black in Figures 9 and 10).

For the 20-cell subset of connection pattern 1, considering only intrinsic spiking histories, we obtained predictive powers in the range of 0.32 to 0.45 with a mean of 0.38; in the full ensemble history case, we obtained values in the range of 0.47 to 0.64 with a mean of 0.57. For the 40-cell subset of connection pattern 1, predictive powers for the intrinsic spiking history case yielded a range of 0.34 to 0.48, with a mean of 0.41, whereas in the full ensemble history case, the predictive power ranged between 0.48 and 0.74, with a mean of 0.64.

For the 20-cell subset of connection pattern 2, considering only intrinsic spiking histories, we obtained predictive powers in the range of 0.26 to 0.56, with a mean of 0.44; in the full ensemble history case, we obtained values in the range of 0.54 to 0.88, with a mean of 0.77. For the 40-cell subset of connection pattern 2, predictive powers for the intrinsic spiking history case yielded a range of 0.08 to 0.57, with a mean of 0.42; whereas in the full ensemble history case, the predictive power ranged between 0.46 and 0.87, with a mean of 0.78. We find that the pattern of connectivity chosen does impact the predictive power.

A natural question that follows is whether one needs the full ensemble of cells to attain the full ensemble predictive power for any particular target cell under consideration. To address this question, we analyzed randomly chosen subsets of 5 cells from the 20- and 40-cell subsets for both connection patterns. For each cell analyzed, we monitored the change in predictive power that occurred as a result of adding a single cell from the remainder of the subset, without replacement, until all extrinsic cells from the subset had been added to the prediction scheme. The particular order in which cells were added was determined by an estimate of the potential influence they would have on the predictive power of the target cell. The results for one cell out of the chosen subset of five, for each subset and from each connection pattern, are shown in Figure 11. We note that a subset of the full ensemble of cells is sufficient to attain a sizable fraction of the maximal predictive power achieved with all cells from the ensemble included. This was borne out for all the cells analyzed.

5 Discussion

The existence of coordinated responses in small ensembles of nearby neurons has received renewed attention, in large part due to a growing number of recent experimental recordings employing a variety of theoretical formalisms (Schneidman et al., 2006; Shlens et al., 2006; Tang et al., 2008; Yu, Huang, Singer, & Nikoli, 2008; Truccolo et al., 2010). We investigated one aspect of this coordination through analyzing the predictive power inherent in the spiking histories of small ensembles of neurons in a computational model of cortex using a point process framework (Truccolo et al., 2005, 2010).

Our main finding is that surprising predictive power exists in small ensembles of modeled cortical neurons (about 20–40 neurons) where these cells are driven solely by random background inputs. A small fraction of the neurons in the model were subject to a synaptic current source large enough to drive an action potential in the neuron with high probability, with the times of synaptic current injection following a Poisson process. This background activity was chosen to mimic random inputs to our section of modeled cortex from external sources. This predictive power points to a strong redundancy in the neural code and highlights the possibility of error-correcting properties at the neuronal level (Schneidman et al., 2006; Truccolo et al., 2010).

We observe interesting effects that manifest due to the connectivity patterns studied. Connectivity pattern 1 is based on a more densely connected network, where layer II/III pyramidal cells generally synapse onto a greater number of layer II/III pyramidal cells in the network (at most 178 for pattern 1 versus 112 for pattern 2). The full ensemble predictive powers attained on average are, however, less for connectivity pattern 1 than those attained

for pattern 2. A partial explanation for this fact is that for the less densely connected network (pattern 2), the pairwise correlation structure is stronger (comparing Figures 2 and 3). In a similar vein, it seems that in the less connected but more strongly correlated network (pattern 2), the addition of fewer extrinsic cells was needed to reach sizable fractions of the full ensemble predictive power (see Figure 11). Although the order in which cells were added was not necessarily the one that maximized the increase in predictive power at each addition, the fact of requiring only a subset of the full ensemble of cells was evident in all cases considered. The small oscillations in the predictive power one observes in Figure 11 once one reaches close to the full ensemble predictive power (in Figures 11b–11d) lie within our estimates for errors in computation of the predictive power ($\sigma \sim < 5\%$) (Hanley & McNeil, 1982). It should be noted that the result that one can reach the full potential predictive power of an ensemble of neurons by observing subgroups of the ensemble has also been alluded to previously (Schneidman et al., 2006; Tkacik, Prentice, Balasubramanian, & Schneidman, 2010).

Our results add to the discussion first pointed out by Truccolo et al. (2010). There, it was shown that small ensembles of neurons in arm-related areas of sensorimotor cortex could reliably predict single-neuron spiking in humans and monkeys performing a range of sensorimotor tasks. They hypothesized that this predictive capability reflected behavior-related encoding, which could not be entirely accounted for by common inputs driving their networks. We show in the light of our results that predictive capabilities lie in locally connected computational networks responding to sparsely distributed random background inputs.

Important differences between our computational model and the recordings outlined in Truccolo et al. (2010) make a direct comparison difficult. We have chosen our ensemble of cells from the same layer of modeled cortex, whereas the results of Truccolo et al. (2010) most likely sample cells from multiple layers. The cross-sectional area from which we have sampled neurons is also likely to have been smaller. From the point of view of network activity, our model displays distinct periods of bursting, which, together with our strategy for choosing our ensembles of neurons, might be responsible for the relatively stronger pairwise correlations we observe. Addressing some of these differences, such as randomizing over layers and sampling neurons spaced farther apart (though of course our model is of limited size), should decrease the magnitude of the substantial predictive power we find, though we doubt this would eliminate its existence altogether. The question of the interpretation of our results is then intimately connected to the (open) question of how to interpret computational models themselves.

That we find predictive power in a separate system not ostensibly driven by strongly correlated inputs is clear. If one accepts that our network and its inputs are not reflective of a piece of motor cortex (for example) performing behavior-related computations, then we have at the very least constructed a realistic system in which significant predictive power exists. Predictive power is then not a unique signature of behavior, though it of course may be implied by it. Furthermore, given the architecture of the computational model studied, our results hint that this predictive power may be a generic property of locally connected networks, and our model then provides an explicit, realistic structure within which one can

begin to unravel the origins of this strongly correlated behavior. An alternative to this view is that our network is reflective of a piece of motor cortex performing behavior-related computations, in which case our results provide further support to those of Truccolo et al. (2010), together with an architecture through which such coordinated activity results.

In this way, we promote the interpretation that locally interconnected networks of neurons contain predictive power such that each unit in the network need not be monitored to obtain knowledge about the collective state of the ensemble. A number of avenues of pursuit present themselves in consequence. As Truccolo et al. (2010) pointed out, the distribution we have constructed does not constitute a minimal model for the system; it simply illustrates that predictive power exists. An interesting question is whether one can do better in terms of prediction through instituting a simpler model, whether it be constructed within the point process framework of Truccolo et al. (2005, 2010), in terms of the maximum entropy approaches presented in the style of Schneidman et al. (2006), Shlens et al. (2006), Tang et al. (2008), Marre, El Boustani, Frégnac, and Destexhe (2009), or indeed through lower-dimensional state-space methods as discussed in Stevenson and Kording (2011). The point process formalism used for prediction was not explicitly tuned to necessarily generate optimal results, and the question of how well such inferential techniques do in terms of extracting available information from spike trains remains an interesting (open) question. In addition, the computational model chosen for analysis was not explicitly tuned to generate the levels of predictive power we observed. How the parameter space of the computational model relates to the observed predictive power is also work for future consideration. The relationship of predictive power to connectivity is another interesting line of inquiry. One imagines there is a relationship between the temporal filters one attains and the connectivity of the network. Whether one can reverse the process and infer connectivity from predictive power would be of significant interest (Okatan et al., 2005; Nykamp, 2007; Kim, Putrino, Ghosh, & Brown, 2011; Chen, Putrino, Ghosh, Barbieri, & Brown, 2011). Indeed the setting of a reduced computational model, where connectivity and network architecture are under direct control, would be a promising place to begin.

Acknowledgments

We thank Piotr Franaszczuk and Sridevi Sarma for discussions. Work at Harvard and at Johns Hopkins was supported in part by the Charles H. Hood Foundation and by the NIH (1K08NS066099-01A1).

References

- Anderson WS, Azhar F, Kudela P, Bergey GK, Franaszczuk PJ. Epileptic seizures from abnormal networks: Why some seizures defy predictability. *Epilepsy Res.* 2012; 99:202–213. [PubMed: 22169211]
- Anderson WS, Kudela P, Cho J, Bergey GK, Franaszczuk PJ. Studies of stimulus parameters for seizure disruption using neural network simulations. *Biol Cybern.* 2007; 97:173–194. [PubMed: 17619199]
- Anderson WS, Kudela P, Weinberg S, Bergey GK, Franaszczuk PJ. Phase-dependent stimulation effects on bursting activity in a neural network cortical simulation. *Epilepsy Res.* 2009; 84:42–55. [PubMed: 19185465]
- Av-Ron E. The role of a transient potassium current in a bursting neuron model. *J Math Biol.* 1994; 33:71–87. [PubMed: 7836871]

- Av-Ron E, Parnas H, Segel LA. A basic biophysical model for bursting neurons. *Biol Cybern.* 1993; 69:87–95. [PubMed: 8334193]
- Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol Cybern.* 1988; 59:189–200. [PubMed: 3179344]
- Brown, EN.; Barbieri, R.; Eden, UT.; Frank, LM. Likelihood methods for neural spike train data analysis. In: Feng, J., editor. *Computational neuroscience: A comprehensive approach*. London: CRC Press; 2004.
- Chen Z, Putrino DF, Ghosh S, Barbieri R, Brown EN. Statistical inference for assessing functional connectivity of neuronal ensembles with sparse spiking data. *IEEE Trans Neural Syst Rehabil Eng.* 2011; 19:121–135. [PubMed: 20937583]
- Chornoboy ES, Schramm LP, Karr AF. Maximum likelihood identification of neural point process systems. *Biol Cybern.* 1988; 59:265–275. [PubMed: 3196770]
- Daley, DJ.; Vere-Jones, D. *An introduction to the theory of point processes*. New York: Springer; 2003.
- Douglas RJ, Martin KAC. Neuronal circuits of the neocortex. *Annu Rev Neurosci.* 2004; 27:419–451. [PubMed: 15217339]
- Fawcett T. An introduction to ROC analysis. *Pattern Recognit Lett.* 2006; 27:861–874.
- Gunning D, Adams C, Cunningham W, Mathieson K, O'Shea V, Smith KM, et al. 30 μ m spacing 519-electrode arrays for in vitro retinal studies. *Nuc Inst Methods A.* 2005; 546:148–153.
- Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology.* 1982; 143:29–36. [PubMed: 7063747]
- Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol.* 1952; 117:500–544. [PubMed: 12991237]
- Kim S, Putrino D, Ghosh S, Brown EN. A Granger causality measure for point process models of ensemble neural spiking activity. *PLoS Comput Biol.* 2011; 7:e1001110. [PubMed: 21455283]
- Kudela P, Franaszczuk PJ, Bergey GK. Changing excitation and inhibition in simulated neural networks: Effects on induced bursting behavior. *Biol Cybern.* 2003; 88:276–285. [PubMed: 12690486]
- Marre O, El Boustani S, Frégnac Y, Destexhe A. Prediction of spatiotemporal patterns of neural activity from pairwise correlations. *Phys Rev Lett.* 2009; 102:138101. [PubMed: 19392405]
- Martignon L, Deco G, Laskey K, Diamond M, Freiwald W, Vaadia E. Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput.* 2000; 12:2621–2653. [PubMed: 11110130]
- Nieuwenhuys R. The neocortex: An overview of its evolutionary development, structural organization and synaptology. *Anat Embryol (Berl).* 1994; 190:307–337. [PubMed: 7840420]
- Nykamp DQ. A mathematical framework for inferring connectivity in probabilistic neuronal networks. *Math Biosci.* 2007; 205:204–251. [PubMed: 17070863]
- Okatan M, Wilson MA, Brown EN. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput.* 2005; 17:1927–1961. [PubMed: 15992486]
- Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *J Neurosci.* 2005; 25:11003–11013. [PubMed: 16306413]
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, et al. Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature.* 2008; 454:995–999. [PubMed: 18650810]
- Schneidman E, Berry MJ II, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature.* 2006; 440:1007–1012. [PubMed: 16625187]
- Segev R, Goodhouse J, Puchalla J, Berry MJ II. Recording spikes from a large fraction of the ganglion cells in a retinal patch. *Nat Neurosci.* 2004; 7:1155–1162.
- Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, et al. The structure of multi-neuron firing patterns in primate retina. *J Neurosci.* 2006; 26:8254–8266. [PubMed: 16899720]

- Stevenson IH, Kording KP. How advances in neural recording affect data analysis. *Nat Neurosci*. 2011; 14:139–142. [PubMed: 21270781]
- Suner S, Fellows MR, Vargas-Irwin C, Nakata GK, Donoghue JP. Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Trans Neural Syst Rehab Eng*. 2005; 13:524–541.
- Tang A, Jackson D, Hobbs J, Chen W, Smith JL, Patel H, et al. A maximum entropy model applied to spatial and temporal correlations from cortical networks *in vitro*. *J Neurosci*. 2008; 28:505–518.
- Tka ik G, Prentice JS, Balasubramanian V, Schneidman E. Optimal population coding by noisy spiking neurons. *PNAS*. 2010; 107:14419–14424. [PubMed: 20660781]
- Truccolo, W. Stochastic models for multivariate neural point processes: collective dynamics and neural decoding. In: Grün, S.; Rotter, S., editors. *Analysis of parallel spike trains*. New York: Springer; 2010.
- Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol*. 2005; 93:1074–1089. [PubMed: 15356183]
- Truccolo W, Hochberg LR, Donoghue JP. Collective dynamics in human and monkey sensorimotor cortex: Predicting single neuron spikes. *Nat Neurosci*. 2010; 13:105–111. [PubMed: 19966837]
- Welsh WF. On the reliability of cross-correlation function lag determinations in active galactic nuclei. *PASP*. 1999; 111:1347–1366.
- Yu S, Huang D, Singer W, Nikoli D. A small world of neural synchrony. *Cereb Cortex*. 2008; 18:2891–2901. [PubMed: 18400792]
- Zhao M, Iyengar S. Nonconvergence in logistic and Poisson models for neural spiking. *Neural Comput*. 2010; 22:1231–1244. [PubMed: 20100077]

Appendix: The Computational Model

The computational model used here constitutes a version of that described in Anderson et al. (2007, 2009) and Anderson, Azhar, Kudela, Bergey, and Franaszczuk (2012). Appendix A of Anderson et al. (2007) contains a more complete description of a reduced version of the model under consideration, though for completeness, we highlight portions of their discussion relevant to the model used here. This earlier version of the model contains the same cell types (as discussed below) and cell classes (e.g., layer II/III pyramidal), arranged in the same columnar structure as for the later version of the model (as featured in Anderson et al., 2009, 2012, and studied in this letter), together with similar intrinsic and extrinsic connectivity schemes. The main difference is the increased number of minicolumns arranged contiguously, taking the total number of cells from 16,384 to 65,536. We note also that the two models referred to above have been posted to the Yale Sense-Lab ModelDB database of computational neuroscience models (accession numbers 98902 and 141507) at <http://senselab.med.yale.edu/modeldb/>.

Each of the 65,536 neurons is modeled in single-compartment format in accord with the excitable membrane model of Hodgkin and Huxley (1952), as modified by Av-Ron et al. (1993) and Av-Ron (1994). The membrane potential V for each cell varies as

$$C \frac{dV(t)}{dt} = I_{syn}(t) - I_{Na}(t) - I_{Ca}(t) - I_K(t) - I_{K(Ca)}(t) - I_A(t) - I_L(t), \quad (A.1)$$

where C is the membrane capacitance; $I_{syn}(t)$ is the input synaptic current; $I_{Na}(t)$ is the inward sodium current; $I_{Ca}(t)$ is the inward calcium current; $I_K(t)$ is the outward potassium

current, which includes a delayed rectifier current; $I_{K(Ca)}(t)$ is a calcium-dependent potassium current; $I_A(t)$ is a transient potassium current; and $I_L(t)$ represents a leak current. The voltage dependence of these currents are presented in the appendix of Anderson et al. (2007). The synaptic current $I_{syn}(t)$ for any neuron is determined through a weighted linear combination over all N_{syn} synapses, where

$$I_{syn}(t) = \sum_{i=1}^{N_{syn}} w_i g_i(t - \tau_i)(V - E_i). \quad (A.2)$$

The sum extends over all N_{syn} synapses onto the neuron, w_i is a weight that modulates the contribution of the i th synapse, and $g_i(t)$ is the conductance function for synapse i , which can depend on a delay τ_i . The term E_i is the reversal potential utilized in the driving function for synapse i . The conductance functions $g_i(t)$ can be expressed as

$$g_i(t) = g_{syn}^i \sum_{j=1}^{N(t)} \left(e^{-\frac{t_j - t}{\tau_d}} - e^{-\frac{t_j - t}{\tau_o}} \right), \quad (A.3)$$

where g_{syn}^i is a conductance constant adjusted to provide each postsynaptic potential with an equivalent amount of injected charge integrated over time. The sum runs over the nonnegligible past action potentials received by synapse i , and τ_d and τ_o are the decay and onset times for a given postsynaptic potential, respectively.

Each neuron was modeled as one of three different types—regular spiking, intrinsic bursting, or fast spiking—depending on its intrinsic electro-physiological properties. Model parameters that characterize each one of these types are presented in Table 2 of Anderson et al. (2007). The layer II/III pyramidal cell ensembles utilized in this paper were of the regular spiking type. Figure 1a of Anderson et al. (2007) displays their response to an internal current injection, which notably exhibits a lack of self-excitation after the source is shut off. Specific parameter values governing their dynamics are presented in the second column of Table 2 in Anderson et al. (2007).

Connectivity between cells is described in detail in sections 2.1 and 2.2 of Anderson et al. (2007) and in Anderson et al. (2009). Most pertinent to our discussion here is the number of layer II/III pyramidal cells contacted by each layer II/III pyramidal cell in our chosen connectivity scheme. By way of example, for connection pattern 1, layer II/III pyramidal cells make isotropic connections outside their minicolumn to a radius of 300 μm , contacting a total of at most 178 layer II/III pyramidal cells along the way. The layer II/III pyramidal cells also contact (extrinsic to their own minicolumn) the layer IV stellate cells, layer V pyramidal cells, basket, and chandelier cells. Anderson et al. (2007) describe (in section 2) a basic connectivity pattern that is then varied by multiplicative factors acting on the total number of connections in the model to reproduce a variety of bursting behaviors with a range of frequencies. In this letter, we have set the relative numbers of inhibitory to excitatory connections to most resemble pattern C (Table 1 of Anderson et al. (2007)). For connection pattern 2, the total number of contacted layer II/III pyramidal cells out to 300 μm

for each layer II/III pyramidal cell was reduced to at most 112. In both patterns, connections were reduced for cells within 300 μm of the edge of the model to help mitigate boundary reflection effects.

Figure 1.

Cell locations of the 20 and 40 randomly selected layer II/III pyramidal cell subsets. The 20-cell subset appears in gray, the 40-cell subset in black. The uppermost layer of the model is schematically displayed, spanning 128 layer II/III pyramidal cells on each side.

Figure 2.

Connection pattern 1 spiking behavior. (a) Raster plot showing spike occurrences for 20 randomly selected layer II/III pyramidal cells from the computational model of Anderson et al. (2007, 2009) during a bursting phase. The burst lasts about 600 ms and spontaneously terminates. (b) Raster plot showing spike occurrences for the 40-cell subset. (c) Correlation coefficients computed for the 20-cell subset and (d) the 40-cell subset. Following Truccolo et al. (2010), they were obtained by considering the extremal Pearson correlation coefficient over all time lags between ± 500 ms.

Figure 3.

Connection pattern 2 spiking behavior. (a) Raster plot showing spike occurrences for 20 randomly selected layer II/III pyramidal cells from the computational model of Anderson et al. (2007, 2009) during a bursting phase. The burst lasts about 800 ms and spontaneously terminates. (b) Raster plot showing spike occurrences for the 40-cell subset. (c) Correlation coefficients computed for the 20-cell subset and (d) the 40-cell subset. Following Truccolo et al. (2010), they were obtained by considering the extremal Pearson correlation coefficient over all time lags between ± 500 ms.

Figure 4.

Raised cosine functions. (a) For the intrinsic history case, $N = 10$ functions were used, onto which the previous 100 ms of spiking history was projected for the target cell under consideration. (b) For the extrinsic history case, $N^+ = 4$ functions were used, onto which the previous 100 ms of spiking history was projected for each external cell considered in the remainder of the ensemble.

Figure 5.

Connection pattern 1 exponentiated temporal filters. (a) Intrinsic history filter for a single cell (cell 12) in the 20-cell subset. The x -axis denotes time prior to the time bin of interest, and the y -axis reflects the change in the baseline probability of firing were a spike to have occurred at x ms in the past. A gain of greater than one denotes an increase in the probability of spiking beyond the baseline probability of spiking, a gain of less than one denotes a decrease in this spiking propensity, and a gain of one denotes no change in this spiking propensity. An effective refractory period is evident, together with an increase in spiking probability, if a spike occurs in the range of about 35 to 60 ms, as well as a longer-term depression of spiking probability based on intrinsic activity from about 60 to 100 ms prior to the time bin of interest. (b) Extrinsic history filters for the remaining 19 cells from the 20-cell subset. The greatest increase in the spiking activity of the target cell (cell 12) is controlled by spikes in the ensemble occurring at about 20 to 30 ms prior to the bin of interest. (c, d) the analogous filters for cell 31 from the 40-cell subset. In each subplot, solid gray lines indicate estimates for uncertainties derived from the standard error associated with parameter fitting.

The insets in panels b and d display uncertainties for a single chosen extrinsic temporal filter.

Figure 6.

Connection pattern 2 exponentiated temporal filters. (a) Intrinsic history filter for a single cell (cell 12) in the 20-cell subset. The x -axis denotes time prior to the time bin of interest, and the y -axis reflects the change in the baseline probability of firing were a spike to have occurred at x ms in the past. A gain of greater than one denotes an increase in the probability of spiking beyond the baseline probability of spiking, a gain of less than one denotes a decrease in this spiking propensity, and a gain of one denotes no change in this spiking propensity. An effective refractory period is evident, together with an increase in spiking probability, if a spike occurs in the range of about 35 to 80 ms. (b) Extrinsic history filters for the remaining 19 cells from the 20-cell subset. The greatest increase in the spiking activity of the target cell (cell 12) is controlled by spikes in the ensemble occurring at about 20 to 40 ms prior to the bin of interest. (c, d) The analogous filters for cell 31 from the 40-cell subset. Solid gray lines represent uncertainty estimates as discussed in the caption to Figure 5.

Figure 7.

Connection pattern 1 conditional probabilities. (a) Predicted conditional probabilities of spiking for target cell 12 from the 20-cell subset, displayed over the same period of time as in Figure 2a. The occurrence of a spike in the target cell is indicated by a vertical gray line. The baseline firing rate is evident in the case where no spikes have occurred in the past 100 ms. One example of the combined intrinsic and ensemble activity indicating the increased likelihood of an upcoming spike in the target neuron is evident within the 20 ms or so prior to the last spike shown on the plot. (b) Predicted conditional probabilities of spiking for target cell 31 from the 40-cell subset displayed over the same period of time as in Figure 2b.

Figure 8.

Connection pattern 2 conditional probabilities. (a) Predicted conditional probabilities of spiking for target cell 12 from the 20-cell subset, displayed over the same period of time as in Figure 3a. The occurrence of a spike in the target cell is indicated by a vertical gray line. The baseline firing rate is evident in the case where no spikes have occurred in the past 100 ms. (b) Predicted conditional probabilities of spiking for target cell 31 from the 40-cell subset displayed over the same period of time as in Figure 3b.

Figure 9.

Connection pattern 1. (a) ROC curves for our prediction scheme for a single cell (cell 12) from the 20-cell subset. (b) ROC curves for a single cell (cell 31) from the 40-cell subset. The gray trace in both cases results from thresholding over conditional probabilities generated by considering only the intrinsic spiking dynamics of the target cell. The black trace shows the case where the external cells are included in the analysis. (c, d) Histograms of the predictive powers obtained for the 20-cell subset (c) and the 40-cell subset (d). The gray bars indicate prediction based solely on the intrinsic spiking history of the target cell, whereas the black bars indicate predictive powers obtained by including the remainder of the ensemble.

Figure 10.

Connection pattern 2. (a) ROC curves for our prediction scheme for a single cell (cell 12) from the 20-cell subset. (b) ROC curves for a single cell (cell 31) from the 40-cell subset. The gray trace in both cases results from thresholding over conditional probabilities generated by considering only the intrinsic spiking dynamics of the target cell. The black trace shows the case where the external cells are included in the analysis. (c, d) Histograms of the predictive powers obtained for the 20-cell subset (c) and the 40-cell subset (d). The gray bars indicate prediction based solely on the intrinsic spiking history of the target cell, whereas the black bars indicate predictive powers obtained by including the remainder of the ensemble.

Figure 11.

Increment in the predictive power as a result of the addition of single cells, without replacement. (a) Cell 13 from connection pattern 1. (b) Cell 13 from connection pattern 2. (c) Cell 24 from connection pattern 1. (d) Cell 24 from connection pattern 2.