Quantifying neuronal network dynamics through coarse-grained event trees

Aaditya V. Rangan, David Cai*, and David W. McLaughlin*

Courant Institute of Mathematical Sciences, New York University, New York, NY 10012

Contributed by David W. McLaughlin, May 2, 2008 (sent for review January 25, 2008)

Animals process information about many stimulus features simultaneously, swiftly (in a few 100 ms), and robustly (even when individual neurons do not themselves respond reliably). When the brain carries, codes, and certainly when it decodes information, it must do so through some coarse-grained projection mechanism. How can a projection retain information about network dynamics that covers multiple features, swiftly and robustly? Here, by a coarse-grained projection to event trees and to the event chains that comprise these trees, we propose a method of characterizing dynamic information of neuronal networks by using a statistical collection of spatial–temporal sequences of relevant physiological observables (such as sequences of spiking multiple neurons). We demonstrate, through idealized point neuron simulations in small networks, that this event tree analysis can reveal, with high reliability, information about multiple stimulus features within short realistic observation times. Then, with a large-scale realistic computational model of V1, we show that coarse-grained event trees contain sufficient information, again over short observation times, for fine discrimination of orientation, with results consistent with recent experimental observation.

information transmission | neuronal coding | orientation selectivity | primary visual cortex

Encoding of sensory information by the brain is fundamental to its operation (1, 2); thus, understanding the mechanisms by which the brain accomplishes this encoding is fundamental to neuroscience. Animals appear to respond to noisy stimulus swiftly within a few 100 ms (3–7). Hence, an immediate important question is what statistical aspects of network dynamics in the brain underlie the robust and reliable extraction of the salient features of noisy input within a short observation window $T_{obs} = \mathcal{O}(100\text{ ms})$ (7, 8). Although the full spatiotemporal history of the high-dimensional network dynamics might contain all of the salient information about the input, an effective and efficient method for extracting the relevant information ultimately entails a projection or "coarse-graining" of the full dynamics to lower dimensions. To be successful, this projection must retain and effectively capture the essential features of noisy input, in a robust and reliable manner, over short observation times T_{obs}. Quantifying neuronal network dynamics by information carried by the firing rates of individual neurons is certainly low-dimensional, but it may require excessively long integration windows when the firing rate is low (9, 10). Here, we propose a method for quantifying neuronal network dynamics by a projection to event trees, which are statistical collections of sequences of spatiotemporally correlated network activity over coarse-grained times. Through idealized networks and through a large-scale realistic model (11, 12) of mammalian primary visual cortex (V1), we show that this event tree-based projection can effectively and efficiently capture essential stimulus-specific, and transient, variations in the full dynamics of neuronal networks. Here, we demonstrate that the information carried by the event tree analysis is sufficient for swift discriminability (i.e., the ability to discriminate, over a short T_{obs}, fine input features, allowing for the reliable and robust discrimination of similar stimuli). We also provide evidence that suggests that, because of their dimensionality, event trees might be capable of encoding many distinct stimulus features simultaneously[†] (note that n features constitute an n-dimensional space that characterizes the input). The idealized networks presented here establish proof of concept for the event tree analysis; the large-scale V1 example presented here indicates that event tree methods might be able to extract fine features coded in real cortices, and our computational methods for analyzing event trees may extend to useful algorithms for experimental data analysis.

Many hypotheses about coding processes in neuronal networks [such as in synfire chains (14)] postulate that the individual spikes or spike rates from specific neurons constitute signals or information packets that can be tracked as they propagate from one neuron to another (15–17). This notion of signal propagation is essentially a feedforward concept; hence, it is restricted to feedforward architecture, where the cascade of signals across neurons in the network can be treated as a causal flow of information through the network (10, 15–17). In contrast, in our event tree analysis, each individual firing event of a particular neuron is never treated as a signal as such. Instead, the entire event tree serves as the signal within the network. Event trees carry information that is a network-distributed (or space–time-distributed) signal, which is a function of both the dynamic regime of the network and its architecture. Here, we will show that this event tree signal can be quantified collectively and statistically without restriction to any particular type of network architecture. In addition, as will be shown below, information represented through the event tree of a network, such as reliability and precision, can differ greatly (and be improved) from those of individual neurons that constitute the components of that network.

Results

To describe and understand the event tree method, it is useful first to recall the information-theoretic framework (1, 2) of "type analysis" (18, 19), a standard projection down to state chains to analyze the dynamics of a system of N-coupled neurons that interact through spikes. Type analysis consists of (*i*) reducing the high-dimensional network dynamics to a raster (a sequence of neuronal firing events); (*ii*) coarse graining time into τ-width bins and recording within each bin the binary "state vector" of spiking neurons within the network (within each time bin, each neuron can spike or not, and a system of N neurons has 2^N possible states); and (*iii*) estimating the conditional probabilities of observing any particular state, given the history of inputs and m previous system states. Type analysis suffers from the curse of dimensionality: it is difficult

Author contributions: A.V.R. and D.C. designed research; A.V.R., D.C., and D.W.M. performed research; and A.V.R., D.C., and D.W.M. wrote the paper.

The authors declare no conflict of interest.

*To whom correspondence may be addressed. E-mail: cai@cims.nyu.edu or david.mclaughlin@nyu.edu.

[†]We note that this feature has been discussed within spike metric coding (13).

This article contains supporting information online at www.pnas.org/cgi/content/full/0804303105/DCSupplemental.

to obtain, over realistically short observation times, accurate statistical approximations to the probabilities of observing any particular sequence of states. For example, if the total observation time of the system $T_{obs} \sim m\tau$, then only $\sim$1 sequence of m states is observed (of a possible 2^{Nm} such sequences). Therefore, this curse limits the ability of type analysis to characterize a short observation of a system.

In contrast, our notion of an event tree invokes a different projection of the system dynamics, namely, down to a set of event chains, instead of state chains. To define event chains, we need the following notation: let σ_t^j denote a firing event of the jth neuron at time t (not discretized), and let σ_I^j denote any firing event of the jth neuron that occurs during the time interval I. Now, given any time scale α, an m-event chain, denoted by $\{\sigma^{j_1} \rightarrow] \; \sigma^{j_2} \rightarrow \ldots \rightarrow \sigma^{j_m}\}$ (spanning the neurons $j_1, \ldots, j_m$, which need not be distinct), is defined to be any event $\sigma_t^{j_m}$ conditioned on (i.e., preceded by) the events $\sigma_{[t-\alpha,t)}^{j_{m-1}}, \cdots, \sigma_{[t-(m-1)\alpha,t-(m-2)\alpha]}^{j_1}$. Unlike type analysis, in which both neuronal firing and nonfiring events affect the probability of observing each state chain (19), our event chain construction limits the relevant observables to firing events only, as motivated by the physiological fact that neurons only directly respond to spikes, with no response to the absence of spikes. Indeed, it seems impossible for the brain itself to respond uniquely to each chain of states consisting of both firing and nonfiring events (e.g., even for a small system of $N = 15$ neurons and a history dependence of $m = 4$ states, the number of possible state chains exceeds the number of cells in a single animal).

Given an observation window T_{obs} of the system, one can record every m-event chain for all m up to some m_{max}. Note that the number of observed one-event chains $\{\sigma^{j_1}\}$ corresponds to the total number of spikes of the j_1th neuron during T_{obs}; the number of observed two-event chains $\{\sigma^{j_1} \rightarrow \sigma^{j_2}\}$ corresponds to the total number of spikes on the j_2th neuron that occur within α ms after a spike of the j_1th neuron; and so forth. We will refer to the full collection of all possible $m \leq m_{max}$ event chains with their occurrence counts as the m_{max}-event tree over T_{obs}.[‡]

Fig. 1 provides a simple example of the event chains produced by a network of coupled integrate-and-fire (I&F) neurons (20). The system is driven by two slightly different stimuli I_1 and I_2. The natural interaction time scale in this system is the synaptic time scale $\alpha \approx 4$ ms, and we record all pairs of events in which the second firing event occurs no later than α ms after the first. Three such two-chains, $\{\sigma^3 \rightarrow \sigma^2\}$, $\{\sigma^3 \rightarrow \sigma^1\}$, and $\{\sigma^2 \rightarrow \sigma^1\}$, are highlighted (within Fig. 1*G*) by light, dark, and medium gray, respectively. Note that the events σ^1, σ^2, σ^3 each occurs two times within both rasters in Fig. 1 *A* and *B*. Fig. 1 *C* and *D* shows representations of the two-event tree corresponding to *A* and *B*, respectively. Note that the event chain $\{\sigma^3 \rightarrow \sigma^1\}$occurs twice within raster *B* but zero times within raster *A*, whereas the event chain $\{\sigma^1 \rightarrow \sigma^3\}$occurs zero times within raster *B* but twice within raster *A*. Fig. 1 *E* and *F* shows representations of the two-event trees associated with very long "$T_{obs} = \infty$" observations of the dynamics under stimuli I_1 and I_2, respectively (where the occurrence counts have been normalized by $T_{obs} \gg 1$ and displayed as rates).

The event tree as described above is a natural intermediate projection of the system dynamics that is lower dimensional than the set of all state chains [dim = $N^{m_{max}}$ in contrast to dim = $(2^N)^{m_{max}}$], but higher dimensional than, say, the firing rate. Nevertheless, there is still a severe undersampling problem associated with analyzing the set of event trees produced by the network over multiple trials of a given T_{obs}. Namely, given multiple T_{obs} trials, each trial will (in general) produce a different event tree, and it is very difficult to estimate accurately the full joint probability distribution (over multiple trials) of the $\sim N^{m_{max}}$ various event chains comprising the event trees. However, we can circumvent this difficulty by considering first the probability distribution (over multiple trials of T_{obs} ms) of the observation count of each event chain individually and then considering the full collection of all of these observation–distributions of event chains (which we will also refer to as an event tree). It is this object that we will use below to assess the discriminability of network dynamics, i.e., how to classify the stimulus based on a T_{obs} sample of the dynamics. In the remainder of the article, the discriminability function is constructed based on standard classification theory (2), by assuming the observation counts of event chains are independent [for details see *Methods* or Fig. S1 in the supporting information (SI) *Appendix*).

It is important to note that event chains are much more appropriate than state chains for this particular method of

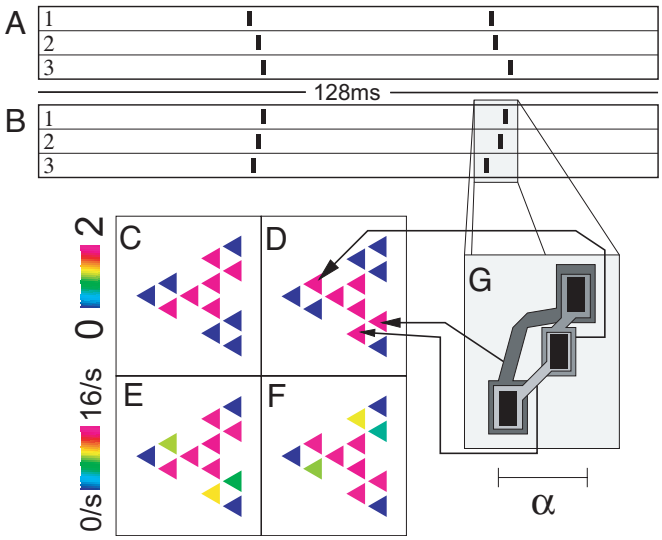

Fig. 1. Illustration of the event chains produced by a network of N coupled I&F neurons (20). For clarity, we choose $N = 3$. The system is driven by two slightly different stimuli, I_1 and I_2. The color scales stand, in general, for the number of occurrences (over T_{obs}) or "occurrence count" of the event $\sigma_t^{j_m}$ *conditioned on* $\sigma_{[t-\alpha,t)}^{j_{m-1}}, \sigma_{[t-2\alpha,t-\alpha)}^{j_{m-2}}, \cdots, \sigma_{[t-(m-1)\alpha,t-(m-2)\alpha)}^{j_1}$. (*A*) A 128-ms raster plot of the network under stimulus I_1. (*B*) Raster plot under stimulus I_2, with the same initial conditions as *A*. (*C* and *D*) Representations of the two-event tree corresponding to *A* and *B*, respectively. The singleton events $\{\sigma^j\}$ of the jth neuron are displayed (within the central triangle) at complex vector location $e^{2\pi i(j-0.5)/3}$ with their occurrence count indicated by color (scale ranging from 0 to 2 recorded events). The occurrence count of event pairs $\{\sigma^j \rightarrow \sigma^k\}$are shown in the peripheral triangles [displayed at complex vector location $3e^{2\pi i(j-0.5)/3} + e^{2\pi i(k-0.5)/3}$]. (*E* and *F*) Representations of the two-event trees associated with very long $T_{obs} = \infty$ observations of the dynamics under stimuli I_1 and I_2, respectively. The color scale stands for the occurrence rate ranging from 0 per s to 16/s. (*G*) This panel zooms in on the second single synchronous burst observed in raster *B*. Within *G* the three black rectangles correspond to spikes, and the three two-chains, $\{\sigma^3 \rightarrow \sigma^2\}$, $\{\sigma^3 \rightarrow \sigma^1\}$, and $\{\sigma^2 \rightarrow \sigma^1\}$, are highlighted (within *G*) by light, dark, and medium gray, respectively. These two-chains each correspond to a different position within the graphic representation of *D*, and these positions are indicated with arrows leading from *G* to *D*. Fig. S1 in the *SI Appendix* further details this graphic representation of event trees, and Fig. S3 in the *SI Appendix* illustrates the utility of this representation.

[‡]An event tree can be thought of as an approximation to the set of conditional probabilities $P_\alpha^{j_1,\cdots,j_m} = P(\sigma_t^{j_m} | \sigma_{[t-\alpha,t)}^{j_{m-1}}, \sigma_{[t-2\alpha,t-\alpha)}^{j_{m-2}}, \cdots, \sigma_{[t-(m-1)\alpha,t-(m-2)\alpha)}^{j_1})$ over the window T_{obs}. Importantly, both T_{obs} and α should be dictated by the dynamics being studied. In many cases, rich network properties can be revealed by choosing T_{obs} comparable to the system memory and α comparable to the characteristic time scale over which one neuron can directly affect the time course of another ($\alpha \approx$2–20 ms). Note that the α-separation of events within each event chain implies that the event tree contains more dynamic information than does a record of event orderings within the network (21).

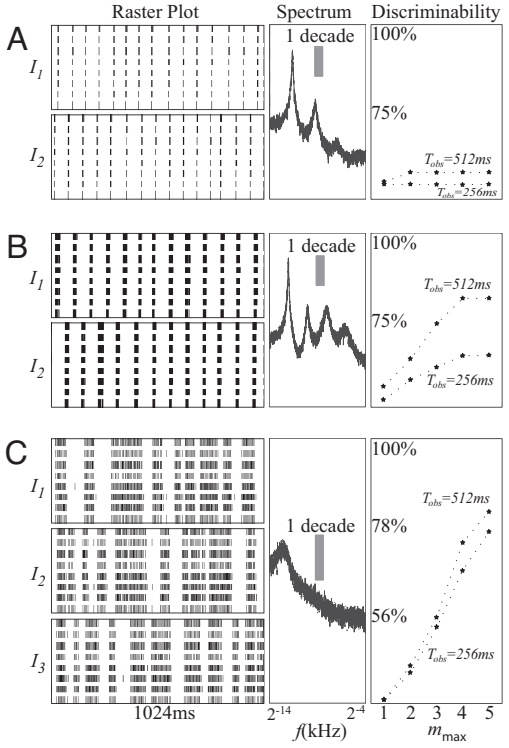

Fig. 2. Swift discriminability. Shown is the utility of event tree analysis through model networks of N conductance-based I&F excitatory and inhibitory neurons, driven by independent Poisson input, for three typical dynamic regimes. For clarity, $N = 8$. The stimuli I_k are fully described by input rate ν_k spikes per ms and input strength f_k, $k = 1, 2, 3$ with $(\nu_1, f_1) = (0.5, 0.005)$, $(\nu_2, f_2) = (0.525, 0.005)$, and $(\nu_3, f_3) = (0.5, 0.00525)$. *A–C* (*Left*) Typical 1024-ms rasters under stimulus I_1, I_2, or I_3. (*Middle*) Log-linear plots of the subthreshold voltage power spectra under stimulus I_1 (dashed) and I_2 (dotted) [and I_3 (dash-dotted) in *C*], which strongly overlap one another. (*Right*) Discriminability of the $T_{obs} = 256$ ms and $T_{obs} = 512$ ms m_{max}-event trees (with $\alpha = 2$ ms) as a function of m_{max}, with the abscissa denoting chance performance. (*A*) Phase oscillator regime. (*B*) Bursty oscillator regime. (*C*) Sustained firing regime. (See *Methods* for details of discriminability).

estimating observation–distributions and assessing discriminability. As we discussed above, there is a curse of dimensionality for state chain analysis: only ~1 sequence of m states is observed over $T_{obs} \sim m\tau$. In contrast, because many distinct event chains can occur simultaneously, there can be a very large number of distinct, stimulus-sensitive event chains (spanning different neurons in the network) even within short ($T_{obs} \sim 100$ ms) observations of networks with low firing rates. Because event chains are not mutually exclusive, multiple event chains can occur during each T_{obs}, and a collection of accurate T_{obs} observation–distributions (one for each event chain) can be estimated with relatively few trials (in contrast to the $O(2^{Nm_{max}})$ trials required to build a collection of observation–distributions of state chains). As will be seen below, it is this statistical feature that enables our event tree projection to characterize robustly, over short T_{obs}, the transient response and relevant dynamic features of a network as a whole (reflecting the dynamic regime the network is in as well as the time-varying inputs). A neuronal network contains information for swift discriminability when that network can generate sufficiently rich, effectively multidimensional event chain dynamics that reflect the salient features of the input, as demonstrated in Figs. 1 and 2. Therefore, we call a network functionally powerful (over T_{obs}) if the event tree (comprising the T_{obs} distribution of event chains) is a sensitive function of the input.[§]

In Fig. 1, the discrepancies between Fig. 1 *C* and *D* are highly indicative of the true discrepancies in the conditional probabilities shown in Fig. 1 *E* and *F*. The $T_{obs} = 128$ ms rasters in Fig. 1 *A* and *B* clearly show that firing rate, oscillation frequency, and type analysis (with $\tau \gtrsim 4$ ms) cannot be used to classify correctly the input underlying these typical $T_{obs} = 128$ ms observations of the system. However, the two-event trees over these $T_{obs} = 128$ ms rasters can correctly classify the inputs (either I_1 or I_2). Furthermore, for this system, a general 128 ms observation is correctly classified by its two-event tree ~85% of the time.

In Fig. 2 we illustrate the utility of event tree analysis for swift discriminability within three model networks (representative of three typical dynamic regimes). The networks are driven by independent Poisson stimuli I_k that are fully described by input rate ν_k spikes per ms and input strength f_k, $k = 1, 2, 3$. The middle panels in Fig. 2 show log-linear plots of the subthreshold voltage power spectra under stimuli I_k. These power spectra strongly overlap one another under different stimuli. With these very similar inputs, the spectral power of synchronous oscillations fails to discriminate the inputs within $T_{obs} \leq 512$ ms. For large changes in the stimulus, these networks can exhibit dynamic changes that are detectable through measurements of firing rate. However, for the cases shown in Fig. 2, with very similar inputs, the firing rate also fails to discriminate the inputs within $T_{obs} \leq 512$ ms.

Fig. 2*A* illustrates a phase oscillator regime, where each neuron participates in a cycle, slowly charging up its voltage under the drive. Every ~70 ms, one neuron fires, pushing many other neurons over threshold to fire, so that every neuron in the system either fires or is inhibited. Then the system starts the cycle again. In this regime, the synchronous network activity strongly reflects the architectural connections but not the input. Note that here the order of neuronal firing within each synchronous activity is independent of the order within the previous one because the variance in the input over the silent epoch is sufficient to destroy the correlation between any neurons resulting from the synchronous activity (data not shown). In this simple dynamic state, neither the firing rate, the power spectrum, nor event tree analysis can reliably discriminate between the two stimuli within $T_{obs} \leq 512$ ms. This simple state, with oscillations, is not rich enough dynamically to discriminate between these stimuli.

Fig. 2*B* illustrates a bursty oscillator regime, where the dynamics exhibits long silent periods punctuated by ~10- to 20-ms synchronous bursts, during which each neuron fires 0–10 times. The power spectrum and firing rates again cannot discriminate the stimuli, whereas deeper event trees ($m_{max} = 4, 5$) here can reliably differentiate I_1 and I_2 within T_{obs} ~512 ms. (As a test of statistical significance, the discriminability computed by using an alternative event tree with neuron labels shuffled across each event chain performs no better than mere firing rates.) We comment that we can also use different time scales α for measuring event trees. For example, in this bursty oscillator regime, we estimated that the variance in input over the silent periods of T_s ~80 ms cannot sufficiently destroy the correlation between neurons induced by the synchronous bursting. Thus, event trees constructed with $\alpha \sim T_s$, observed across silent periods by including multiple sustained bursts, can also be used to discriminate the inputs (data not shown).

Fig. 2*C* illustrates a sustained firing regime, where the power spectrum and one-event tree (i.e., firing rates) cannot discrim-

[§]The T_{obs} distribution of event chains is a statistical collection that includes the occurrence counts of every α-admissible event chain, and is more sensitive to stimulus than the rank order of neuronal firing events (21).

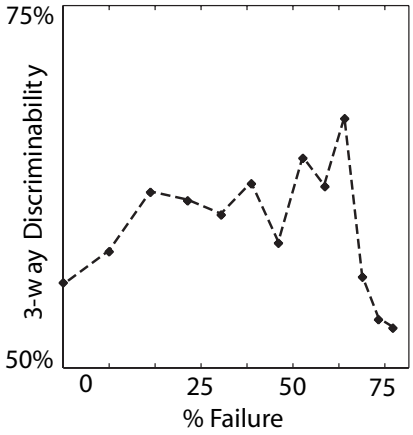

Fig. 3. Functional power of the network vs. reliability of individual neurons. Here, we illustrate that the functional power of a network is not simply an increasing function of the synaptic reliability of individual neurons within that network. More details about this model network may be found in Fig. S2 in the *SI Appendix*. The discriminability is plotted for a set of different p_{fail} (the dots indicate data points, the dashed line is to guide the eye only).

inate the stimuli well (chance performance for the discrimination task is 33%), whereas the deeper event chains can discriminate between the multiple stimuli very well. The five-event tree over $T_{\text{obs}} \lesssim 256$ ms can be used to classify correctly the stimulus ~75% of the time. Incidentally, a label-shuffled event tree performs the discriminability task at nearly chance (i.e., firing rate) level. The fact that the five-event tree can be used to distinguish among these three stimuli implies that event tree analysis could be used to discriminate robustly between multiple stimuli (such as f and v).

To summarize, the dynamics shown in Fig. 2 *B* and *C* is sufficiently rich that the event trees observed over a short $T_{\text{obs}} \lesssim 256$ ms can (*i*) reliably encode small differences in the stimulus and (*ii*) potentially serve to encode multiple stimuli as indicated in Fig. 2*C*.

In the present work we did not investigate the map from high-dimensional stimulus space to the space of observation–distributions of event chains (13). However, we have tested the ability of the sustained firing regime (see Fig. 2*C*) to distinguish between up to six different stimuli (which differ along different stimulus dimensions) simultaneously. We chose uniform independent Poisson stimuli I_j such that: (*i*) I_1 had fixed strength f and rate v; (*ii*) I_2 had strength f_2 and rate v; (*iii*) I_3 had strength f and rate v_3; (*iv*) I_4 had strength f and rate $v_4[\cos(wt)]_+$, a rectified sinusoid oscillating at 64 Hz; (*v*) I_5 had strength f and rate $v_5[\cos(2wt)]_+$, a rectified sinusoid oscillating at 128 Hz; and (*vi*) I_6 had strength f and rate given by a square wave oscillating at 64 Hz and amplitude v_6. We fixed f_2, v_3, v_4, v_5, v_6 so that the firing rates observed under stimuli $I_1, \ldots, I_5$ were approximately the same. Specifically, within this six-stimulus discrimination task, the one-, two-, three-, four-, and five-event trees over $T_{\text{obs}} \lesssim 512$ ms could be used to classify correctly the stimulus ~18%, ~20%, ~22%, ~25%, and ~34% of the time, respectively. Only the deeper event trees contained sufficient information over T_{obs} to discriminate the stimuli at a rate significantly greater than the chance level of 17%. Again, for this discrimination task, label-shuffled event trees perform at nearly chance (i.e., firing rate) level.

We emphasize that the functional power of a network is not simply related to the individual properties of the neurons composing the network. For instance, the functional power of a network can increase as its components become less reliable, as is illustrated in Fig. 3 (and described in more detail in Fig. S2 in the *SI Appendix*). Fig. 3 shows an example of a model network whose discriminability increases as the probability of synaptic failure (20) increases, making the individual synaptic components of the network less reliable. Here, we model synaptic failure by randomly determining whether each neuron in the network is affected by any presynaptic firing event with $p_{\text{trans}} = (1 - p_{\text{fail}})$ as transmission probability. We estimate the functional power of this network as a function of p_{fail}. The system was driven by three similar inputs, I_1, I_2, I_3, and we record the $T_{\text{obs}} = 512$ ms three-event trees. We use the three-event trees to perform a three-way discrimination task (33% would be chance level). The discriminability is plotted against p_{fail} in Fig. 3, which clearly demonstrates that the event trees associated with the network are more capable of fine discrimination, when $p_{\text{fail}} \sim 60\%$ than when $p_{\text{fail}} = 0\%$. If there is no synaptic failure in the network, then the strong recurrent connectivity within the network forces the system into one of two locked states. However, the incorporation of synaptic failure within the network allows for richer dynamics. A possible underlying mechanism for this enhanced reliability of the network is "intermittent desuppression" (12): synaptic failure may "dislodge" otherwise locked, input-insensitive, responses of the system. As a consequence, the dynamics escapes from either of these locked states and generates more diverse input-sensitive event trees over a short T_{obs}, thus leading to a system with a higher sensitivity to inputs and a higher functional power.

It is important to emphasize that the analysis of network dynamics using information represented in event trees and characterization of functional power can be extended to investigate much larger, more realistic neuronal systems, such as the mammalian V1. Neurons within V1 are sensitive to the orientation of edges in visual space (22). Recent experiments indicate that the correlations among spikes within some neuronal ensembles in V1 contain more information about the orientation of the visual signal than do mere firing rates (23, 24). We investigate this phenomenon within a large scale model of V1 (see refs. 11 and 12 for details of the network).

For these larger networks, it is useful to generalize the notion of events from the spikes of individual neurons to spatiotemporally coarse-grained regional activity, as illustrated in Fig. 4, in which a regional event tree is constructed by using regional events, defined to be any rapid sequence of neuronal spikes (viewed collectively as a single "recruitment event") occurring within one of the N_r cortical regions of the model V1 cortex. More specifically, we define N_r sets of neurons (i.e., "regions") $\{J_i\}$, each composed of either the excitatory or inhibitory neurons within a localized spatial region of the V1 model network (see shaded regions in Fig. 4), and we say that a regional event $\sigma_t^{J_i}$ takes place any time N_{local} neurons within a given region fire within τ_{local} ms of one another. We define the time of the regional event $\sigma_t^{J_i}$ as the time of the final local firing event in this short series. This characterization of regional events within regions of excitatory neurons (using $N_{\text{local}} \sim 3$–5 and $\tau_{\text{local}} \sim 3$–6 ms) serves to quantify the excitatory recruitment events we have observed and heuristically described within the intermittently desuppressed cortical operating point of our cortical V1 model (11). The choice of τ_{local} corresponds to the local correlation time scale in the system, and the choice of N_{local} corresponds to the typical number of neurons involved in recruitment. We have observed (12) that recruitment events in neighboring regions are correlated over time scales of 10–30 ms. These recruitment events are critical for the dynamics of our V1 model network, and the dynamic interplay between recruitment events occurring at different orientation domains can be captured by a regional event tree defined by using $\alpha \sim 15$ ms.

In Fig. 4, we demonstrate that the coarse-grained event tree associated with the dynamics of our large-scale ($\sim 10^6$ neurons) computational, recurrent V1 model (11, 12) is indeed sufficiently rich to contain reliably information for small changes in the input

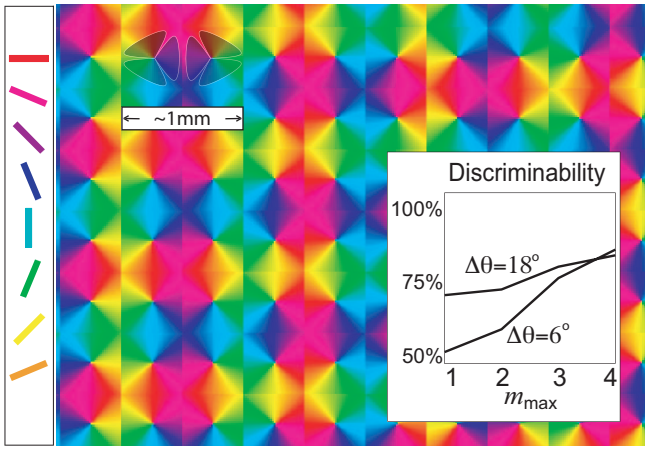

Fig. 4. Generalization of event trees and discriminability of fine orientations in our large-scale numerical V1 model. The colored lattice illustrates the spatial scale and orientation hypercolumns in our V1 model (11, 12). Cells are color-labeled by their preferred orientation shown on the left. To estimate orientation discrimination, we measure a regional event tree based on $N_r = 12$ total regions, comprising the excitatory and inhibitory subpopulations within the six shaded isoorientation sectors, and an α of 15 ms (regional events are defined by using $N_{local} = 3$ and $\tau_{local} = 4$ ms; see *Results*). (*Inset*) Orientation discriminability of the $T_{obs} = 365$ ms regional event tree as a function of m_{max} (*i*) for two similar stimuli ($\Delta\theta = 6°$, lower curve) and (*ii*) for two disparate stimuli ($\Delta\theta = 18°$, upper curve). In both cases, the discriminability computed by using a label-shuffled regional event tree was no better than that computed by using mere regional event rates. In addition, similar conclusions hold for shorter $T_{obs} = 150–250$ ms (data not shown).

orientation. If two stimuli are sufficiently different in terms of angle, then the firing rate alone is sufficient for discrimination. However, deeper event trees are necessary to discriminate between two very similar stimuli (say, two gratings differing by $\lesssim 6°$), slight differences in relative input induce different spatiotemporal orderings of activity events, giving rise to different event trees. Therefore, the observed deeper ($m = 2, 3, 4$) regional event trees can successfully distinguish between the stimuli, whereas the regional activation count ($m = 1$) cannot, as is confirmed in Fig. 4 (*Inset*) and consistent with the experiment (23, 24). This illustrates the possible importance of correlations within neuronal activity. Note that regional event tree analysis using $N_{local} = 1$ (i.e., chains of single-neuron firing events) does not effectively discriminate similar stimuli (data not shown). We comment that the process of pooling (9) regional event trees recorded over multiple different spatial regions to improve discriminability is much more efficient than pooling mere regional event rates (data not shown). Our V1 results suggest that it might be possible to analyze event tree dynamics to reveal information for fine orientations in the real V1.

Discussion

We have proposed a method of quantifying neuronal network dynamics by using information present in event trees, which involves a coarse-grained statistical collection of spatiotemporal sequences of physiological observables. We have demonstrated that these spatiotemporal event sequences (event chains) can potentially provide a natural representation of information in neuronal network dynamics. In particular, event tree analysis of our large-scale model of the primary visual cortex (V1) is shown to provide fine discriminability within the model and hence possibly within V1. Importantly, the event tree analysis is shown to be able to extract, with high reliability, the information contained in event trees that simultaneously encodes various stimuli within realistic short observation times.

The event tree analysis does not rely on specific architectural assumptions such as the feedforward assumption underlying many coding descriptions of synfire chains; in fact, event tree analysis is applicable to both feedforward and strongly recurrent networks. Discriminability relies, particularly for fine discrimination tasks, on the network operating with sufficiently rich dynamics. In this regard, we have demonstrated that networks locked in simple dynamics, no matter whether characterized by oscillation frequencies through power spectrum or event trees, cannot discriminate between fine stimulus characteristics. However, as has been demonstrated above, there exist networks that exhibit complex dynamics that contains sufficient information to discriminate stimuli swiftly and robustly, information that can be revealed through an event tree projection but not by merely analyzing power spectrum of oscillations. Moreover, we have shown that event trees of a network can reliably capture relevant information even when the individual neurons that comprise the network are not reliable. There are other useful high-dimensional projections of network dynamics, such as spike metric (13, 25) and ensemble encoding (26), which may also be capable of extracting information that can be used to discriminate stimuli robustly and swiftly. It would be an interesting theoretical endeavor to investigate these issues by using these alternative projections and to compare their performance with event tree analysis.

We expect that our computational methods for collecting, storing, and analyzing event trees can be used by experimentalists to study network mechanisms underlying biological functions by probing the relevance and stimulus specificity of diverse subsets of events within real networks through methods such as multielectrode grids.

The features described above make the event tree analysis intriguing. Here, we have addressed how information could be represented in a network dynamics through coarse-grained event trees. As an extended space–time projection, it will probably require an extended space–time mechanism for other neurons to read out the information contained in event trees. A theoretical possibility would be that read-out neurons employ the mechanism of tempotron (27) to reveal the information that is represented through event trees.

Finally, we mention a possible analytical representation of the dynamics of event trees in the reduction of network dynamics to a much lower dimensional effective dynamics. For example, in the phase oscillator dynamics of Fig. 2*A*, the event tree observed during one synchronous activity is uncorrelated from that observed during the next synchronous activity. This independence allows us to reduce the original dynamics to a Markov process of successive event trees (data not shown). However, for more complicated dynamics, reduction cannot be achieved by this simple Markov decomposition. Instead, a hierarchy of event chains, namely, chains of chains, needs to be constructed for investigating correlated dynamics over multiple time scales.

Methods

Standard computational model networks of conductance-based I&F point neurons (20), driven by independent Poisson input, are used to test event tree analysis. For details, see *SI Appendix*. For application to V1, we use the large-scale, realistic computational model of conductance based point neurons described in refs. 12 and 28. We emphasize that (*i*) the general phenomena of robust event tree discriminability and (*ii*) the ability of event tree analysis to distinguish between many stimuli that differ along distinct stimulus dimensions are not sensitive to model details and persist for a large class of parameter values and for a wide variety of dynamical regimes.

In practice, one usually cannot estimate the full multidimensional T_{obs} distribution of the m-event tree for a system because the dimension is just too high to estimate effectively such a joint probability distribution of observation counts of all different event chains. To circumvent this curse of dimensionality, we separately estimate the T_{obs} distribution of each m-event chain within the event tree. In other words, we first record many independent samples of this

network's m_{max}-event tree (over multiple independent observations with fixed T_{obs}) under each stimulus I_1, I_2. With this collection of data, we obtain, for each stimulus, the empirical distributions of each m-event chain's occurrence count for all $m \leq m_{max}$. Thus, for each stimulus I_l and for each event chain $\{\sigma^{j1} \rightarrow \sigma^{j2} \rightarrow \ldots \rightarrow \sigma^{jm}\}$ we obtain a set of separate probabilities $P_l\{\sigma^{j1} \rightarrow \sigma^{j2} \rightarrow \ldots \rightarrow \sigma^{jm}\}$ for each chain to occur k times within a given T_{obs}), for each integer $k \geq 0$ and each stimulus $l = 1, 2$. We then apply standard methods from classification theory and use this set of observation count distributions, along with an assumption of independence for observation counts of different event chains, to perform signal discrimination from a single T_{obs} observation. For completeness, we describe our procedure below.

Typically, some event chains are not indicators of the stimulus (i.e., the T_{obs} distribution of occurrence count is very similar for distinct stimuli). However, other event chains are good indicators and can be used to discriminate between stimuli. For example, as depicted in Fig. S1*E* in the *SI Appendix* the $T_{obs} = 512$ ms distribution of occurrence counts of the four-event chain $\sigma^4 \rightarrow \sigma^1 \rightarrow \sigma^2 \rightarrow \sigma^8$ is quite different under stimulus I_1 than under stimulus I_2. Given (estimates of) these two distributions $P_1(\cdot)$ and $P_2(\cdot)$, one obtains from a single T_{obs} measurement the occurrence count p of the $\sigma^4 \rightarrow \sigma^1 \rightarrow \sigma^2 \rightarrow \sigma^8$ event chain. We choose I_1 if $P_1(p) > P_2(p)$, otherwise we choose I_2. Then, we use the two distributions to estimate the probability that this choice is correct, resulting in a hit rate $A = \frac{1}{2}\int_0^\infty \max(P_1, P_2)\, dn$ and a false alarm rate $B = 1 - A$, and the information ratio $I_{I_1, I_2}^{\sigma^{j1} \rightarrow \cdots \rightarrow \sigma^{jm}} \equiv A/B$.

The procedure described above classifies the stimulus underlying a single T_{obs} observation by considering only a single-event chain (i.e., a single element of the event tree). We can easily extend this procedure to incorporate every event chain within the event tree constructed from one T_{obs} observation. For example, given a T_{obs} observation and its associated event tree, we can use the procedure outlined above to estimate, for each chain separately, which stimulus induced the tree. Thus, each event chain ''votes'' for either stimulus I_1 or I_2, weighting each vote with the log of the information ratio $(I_{I_1, I_2}^{\sigma^{j1} \rightarrow \cdots \rightarrow \sigma^{jm}})$. We then sum up the weighted votes across the entire event tree to determine the candidate stimulus underlying the sample T_{obs} observation. We define the discriminability of the m_{max}-event tree (for this two-way discriminability task) to be the percentage of sample observations that were correctly classified under our voting procedure. To perform three-way discriminability tasks, we go through an analogous procedure, performing all three pairwise discriminability tasks for each sample observation and ultimately selecting the candidate stimulus corresponding to the majority. Note that the discriminability is a function of α, T_{obs}, and m_{max}. For most of the systems we have observed, the discriminability increases as m_{max} and T_{obs} increase.

ACKNOWLEDGMENTS. The work was supported by National Science Foundation Grant DMS-0506396 and by a grant from the Swartz Foundation.

1. Rieke F, Warland D, Ruyter van Steveninck RR, Bialek W (1996) *Spikes*: *Exploring the Neural Code* (MIT Press, Cambridge, MA).
2. Dayan P, Abbott LF (2001) *Theoretical Neuroscience* (MIT Press, Cambridge, MA).
3. Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489.
4. Uchida N, Mainen ZF (2003) Speed and accuracy of olfactory discrimination in the rat. *Nat Neurosci* 6:1224–1229.
5. Rousselet GA, Mace MJ, Fabre-Thorpe M (2003) Is it an animal? Is it a human face? Fast processing in upright and inverted natural scenes. *J Vis* 3:440–455.
6. Abraham NM, *et al*. (2004) Maintaining accuracy at the expense of speed: Stimulus similarity defines odor discrimination time in mice. *Neuron* 44:865–876.
7. Mainen ZF (2006) Behavioral analysis of olfactory coding and computation in rodents. *Curr Opin Neurobiol* 16:429–434.
8. Uchida N, Kepecs A, Mainen ZF (2006) Seeing at a glance, smelling in a whiff: Rapid forms of perceptual decision making. *Nat Rev Neurosci* 7:485–491.
9. Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
10. Litvak V, Sompolinsky I, Segev I, Abeles M (2003) On the transmission of rate code in long feedforward networks with excitatory–inhibitory balance. *J Neurosci* 23:3006–3015.
11. Cai D, Rangan AV, McLaughlin DW (2005) Architectural and synaptic mechanisms underlying coherent spontaneous activity in V1. *Proc Natl Acad Sci USA* 102:5868–5873.
12. Rangan AV, Cai D, McLaughlin DW (2005) Modeling the spatiotemporal cortical activity associated with the line–motion illusion in primary visual cortex. *Proc Natl Acad Sci USA* 102:18793–18800.
13. Victor JD, Purpura KP (1997) Sensory coding in cortical neurons: Recent results and speculations. *Ann NY Acad Sci* 835:330–352.
14. Abeles M (1991) *Corticonics*: *Neural Circuits of the Cerebral Cortex* (Cambridge Univ Press, New York).
15. Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
16. van Rossum MC, Turrigiano GG, Nelson SB (2002) Fast propagation of firing rates through layered networks of noisy neurons. *J Neurosci* 22:1956–1966.
17. Vogels TP, Abbott LF (2005) Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci* 25:10786–10795.
18. Johnson DH, Gruner CM, Baggerly K, Seshagiri C (2001) Information–theoretic analysis of neural coding. *J Comput Neurosci* 10:47–69.
19. Samonds JM, Bonds AB (2005) Cooperative and temporally structured information in the visual cortex. *Signal Proc* 85:2124–2136.
20. Koch C (1999) *Biophysics of Computation* (Oxford Univ Press, Oxford, UK).
21. Thorpe SJ, Gautrais J (1998) Rank order coding. *Computational Neuroscience*: *Trends in Research*, ed Bower J (Plenum, New York), pp 113–119.
22. De Valois R, De Valois K (1988) *Spatial Vision* (Oxford Univ Press, New York).
23. Samonds JM, Allison JD, Brown HA, Bonds AB (2003) Cooperation between area 17 neuron pairs enhances fine discrimination of orientation. *J Neurosci* 23:2416–2425.
24. Samonds JM, Allison JD, Brown HA, Bonds AB (2004) Cooperative synchronized assemblies enhance orientation discrimination. *Proc Natl Acad Sci USA* 101:6722–6727.
25. Victor JD (2005) Spike train metrics. *Curr Opin Neurobiol* 15:585–592.
26. Harris KD (2005) Neural signatures of cell assembly organization. *Nat Rev Neurosci* 6:399–407.
27. Gutig R, Sompolinsky H (2006) The tempotron: A neuron that learns spike timing-based decisions. *Nat Neurosci* 9:420–428.
28. Rangan AV, Cai D (2007) Fast numerical methods for simulating large-scale integrate-and-fire neuronal networks. *J Comp Neurosci* 22:81–100.