

On the temporal resolution of neural activity

Akira Date^{*}

Division of Applied Mathematics
Brown University
Providence, RI 02912
date@dam.brown.edu

Elie Bienenstock

Division of Applied Mathematics
Brown University
Providence, RI 02912
elie@dam.brown.edu

Stuart Geman

Division of Applied Mathematics
Brown University
Providence, RI 02912
geman@dam.brown.edu

May 1998

Abstract

An important issue regarding brain function is the existence and role of fine temporal structure in neural activity. Multi-neuronal recording techniques are now available to study this issue. We present a simple statistical method devised to detect fine temporal structure in simultaneously recorded spike processes. We apply this method to data recorded from monkey Supplementary Motor Area, and show a preliminary result which suggests that the nervous system may indeed use a temporal resolution of about 10 ms (or higher).

1 INTRODUCTION

There is an ongoing debate about neural coding: according to some, information is carried only in the average spike frequency, while others hold that the brain may actually

^{*}Department of Biophysical Engineering, Osaka University, Toyonaka, Osaka, 560-8531

use the fine temporal structure of neural discharges. This may include various temporal relationships between spikes emitted by different neurons. Such a view stems in particular from the consideration that the brain needs to somehow bind entities and events represented hierarchically in different cortical areas (Damasio, 1989; Livingston, 1996; von der Malsburg, 1995). If there exists a mechanism for rapidly and reversibly binding otherwise uncorrelated spatio-temporal patterns of neural activities, one could look for evidence of this binding mechanism in the fine temporal structure of single- and multi-neuronal spiking processes.

A related, albeit more general, question is: What is the time scale used by the nervous system? By this we mean the following. Suppose we could randomly and *independently* perturb the time of occurrence of each spike of each neuron by a small amount; would this affect the behavior of the subject? If the perturbation is smaller than “the time scale of the brain,” nothing would change. If on the other hand the perturbation is large enough, we expect that the subject’s behavior will be affected.

Since there is no way to move spikes randomly in real brains, we address this issue indirectly. We ask simply whether there is *any reliably reproducible* fine temporal structure in brains. The alternative, which will be our null hypothesis H_0 , is that brain activity is completely random when looked at on a fine resolution. That is, the multi-dimensional spiking process of any collection of neurons cannot be distinguished from a randomly generated process which has the same properties as the original one on a coarse time scale. Under the null hypothesis, a small perturbation couldn’t possibly affect the functioning of the brain in any way.

To formulate H_0 , we use the following model. Suppose the multidimensional spiking process is an inhomogeneous Poisson process in which all firing rates remain constant in time intervals of length Δ . (We use a fixed, arbitrary, partition of the time axis into intervals of length Δ .) Other than that, the firing rates are themselves a multi-dimensional stochastic process, with any possible joint probability distribution. Depending on Δ , this model can be made completely general: if $\Delta = 1$ ms, then the model includes *all* possible multi-dimensional spiking processes defined on the 1-ms time scale. If however, say, $\Delta = 10$ ms, then the model makes a statement about the randomness of spiking on the fine, i.e., 1-ms, time scale.

To reject the null with $\Delta = 10$, we devise an appropriate *statistical test*. Based on an observed spike sequence from a multi-unit recording, we randomly generate many spike sequences with the same spike count, for each neuron and in each interval of length Δ , as the corresponding spike count in the original sequence. We then ask whether the original observation differs, in *any* statistically significant way, from the jittered ones. The null hypothesis, H_0 , says that the original multi-dimensional point process actually *comes from* the random distribution we use to generate all these jittered processes. To reject this null, we use a *statistic* S , which can be *any* function of the process. If S computed on the original sequence is significantly different from S computed on the jittered data, i.e., from the sample statistic, we reject the null and conclude that there is fine temporal structure in the brain on a time scale that is at least as fine as Δ .

Of course such a conclusion doesn’t imply that the fine time structure is actually *used* by the brain or that this structure is important in any way. Such structure could be an epiphenomenon of the dynamics of interconnected neurons. Also, the mere existence

of a refractory period will in some cases allow us to reject H_0 for the spiking activity of a *single* neuron. Alternatively, we may reject the null because the activities of the recorded neurons are all time-locked to an external event, say a stimulus. Yet in some other cases, the *form* of the statistic S will make it *unlikely* that the structure uncovered is an epiphenomenon. Further, it may suggest that one look for a possible *covariation* of S with some aspect of the animal’s behavior. If such a covariation were shown to exist and to be reliable, this would strongly indicate that fine-temporal structure is used by the brain—in the same way as the covariation of firing rates of V1 cells with stimulus orientation strongly indicates that firing rates are used by the brain.

The statistic S could be chosen from an immense—virtually infinite—collection of functions of a multi-dimensional spike process. Our statistic should have (i) strong power to detect reliably occurring fine temporal structure, and (ii) a possible explanation in terms of a sound neural mechanism. In this study, we looked at statistics that are related to the repeated occurrence of specific spatio-temporal patterns, called “Abeles patterns,” or “synfire events,” which have been reported in the literature (Abeles et al., 1993). We thus ask whether specific spatio-temporal patterns, belonging to a given parametric class of patterns, repeat significantly more often in the original spike sequence than in the jittered ones. For instance, S may count the number of patterns, in that given class, which recur at least 12 times in a record of a given length.

2 A PROBABILITY MODEL FOR SPIKE TRAINS

The multi-spike record is first partitioned into equal-sized bins. We have experimented with bin widths between one and three milliseconds. The “spike process” of the i -th neuron ($s_i(t)$) is then the zero-one process which has value one if there are one or more spikes in the t -th bin, and zero otherwise. Note that time t has been discretized by the bin size $|\text{bin}|$, so that the t -th bin corresponds to times between $[|\text{bin}| \times t]$ ms and $[|\text{bin}| \times t + |\text{bin}| - 1]$ ms in the original record.

Let $p_i(t) = \text{Prob}\{\text{neuron } i \text{ fires at time } t\}$, be the “rate function” of the i -th neuron, and assume that the firing of neuron i at time t is independent of the firing of neuron i at all other times $\tau \neq t$. Note that *any* realization of a spike process can be modeled in this way, if only by taking $p_i(t)$ to be zero or one at each t . In general, $p_i(t)$ can itself be stochastic, to reflect dependence on the task, the surrounding neurons, the time of the recording, and so on.

We define a “time scale” by specifying an interval Δ (which will always be a multiple of the bin size), and we consider the hypothesis that the brain operates on a time scale no smaller than Δ . In other words, we hypothesize that $p_i(t)$ does not vary, or varies slowly, over time intervals of size Δ . We have examined time scales between two and thirty ms.

If $p_i(t)$ does not change appreciably over intervals of size Δ , then perturbing (or “jittering”) individual spikes by as much as Δ ms, should not change in any way the character of the spike train record. To get at this idea more formally, first fix a specific neuron “ i ” (the generalization to a population of neurons will be transparent), and let I_1 be the number of “events” (i.e. number of times $s_i(t) = 1$) for bins contained in the first Δ interval, let I_2 be the number of events in next (disjoint) Δ interval, and so-on. We can now generate, randomly, a large number of (conditionally) independent spike trains,




Figure 1: Random spike movements (spike jittering), a , original spike record b,c , jittered spike records.

T^1, T^2, \dots , from a probability distribution that is essentially (up to the “varies no faster than Δ ” hypothesis) the probability distribution of the original spike train (See Fig. 1): let $T^j, j = 1, 2, \dots$ describe the j 'th sample, which consists of events in bins $t_k^j, k = 0, 1, \dots$. In each T^j the number of events in each interval, I_1, I_2, \dots , is kept constant and equal to the number of events in the original record. Under this constraint, the *locations* of the events are chosen randomly (uniformly) and independently for each interval in each record. Under the hypothesis, $T^i, i = 1, 2, \dots$, are independent and essentially indistinguishable from the original event sequence T , conditioned on I_1, I_2, \dots .

How big is Δ in real brains? In what follows, we will look at a *statistic* $S(T)$ of a multi-spike train T , and ask whether the observed value is significantly different from the values $S(T^i), i = 1, 2, \dots$ calculated from the Monte Carlo sequences. A significant difference rejects the hypothesis that firing rates do not change over Δ -sized intervals.

3 ANALYSIS OF MULTI-ELECTRODE DATA

As remarked earlier, there is essentially an infinite class of statistics that could be employed. For *any* statistic we can choose an appropriate threshold and establish a “significance level” for rejecting the “null” (slow variation in spike rate) hypothesis. However, depending upon the *alternative* hypothesis, such a statistic may or may not have substantial *power*: The probability of rejecting the null hypothesis, given that a specific alternative is in fact true, will depend critically upon the statistic used. It is possible that real spike trains contain, or are in fact the superposition of, more-or-less precise spatio-temporal patterns of activities. In these studies, we choose a statistic that can be reasonably expected to have good power for this class of alternative hypotheses.

More specifically, we will examine the maximum number of repetitions of spatio-temporal patterns of a given “complexity”. A complexity-5 pattern is a spatio-temporal pattern consisting of 5 spike events among the neuronal records. For example, suppose

neuron 2 fires, and after 3 ms neurons 1 and 3 fire simultaneously, and after an additional 30 ms neuron 1 again fires, and, finally, after 50 more ms neuron 4 fires. We call this a spatio-temporal pattern of complexity 5; it is conveniently and succinctly identified by the expression (2,1,3,1,4; 3,0,30,50). In our experiments, patterns are not necessarily resolved at one ms resolution. Instead, they are defined in terms of the *spike processes* $s_i(t)$ introduced earlier. Consequently, spatio-temporal patterns are defined only up to the chosen bin width—in effect, a broader definition of “pattern”.

In general, then, a complexity- k pattern is defined as $(N_1, N_2, \dots, N_k; d_1, d_2, \dots, d_{k-1})$, where each $N_i, i = 1, \dots, k$ specifies a neuron and each $d_i, i = 1, \dots, k - 1$ specifies a delay (inter-spike) interval. Finally, we impose an *a priori* limit on pattern widths by requiring that $\sum_{i=1}^{i=k-1} d_i \leq w$. In our experiments, w ranged from 100 to 500 ms.

Obviously, there are a lot of potential patterns. There is, therefore, the practical problem of finding all instances of repeated patterns of a given complexity. We have developed an algorithm for computing the number of times each complexity- k pattern appears in a multi-record. (See also Abeles et al., 1993, for another approach to a similar computational problem.) In brief outline, the algorithm proceeds as follows: Having fixed a complexity (e.g. 5) and a window width w (e.g. 200 ms), we first detect all pairs of times $(t1, t2)$ at which any pattern of the targeted complexity repeats, where the “time” of a pattern is taken to be the time of its first event. Based upon this pair list, we then develop a list of triplets $(t1, t2, t3)$ —for each triplet there is some pattern that repeats at least three times. The process is continued, recursively, until no patterns are found that repeat more than a given number of times.

We have experimented with spike-train recordings from five neurons, with each record having 1095966 entries of 0’s (no spike) and 1’s (spike), captured at a 1 ms resolution. The recordings were obtained from awake monkeys performing visual-motor tasks over the course of approximately twenty minutes (See Discussion section below).

As a specific example, let us fix the bin width at 3 ms, the maximum pattern width at 100 ms, and the complexity at 4. Table 1 shows the number of patterns of complexity at least 4 that repeat exactly k times, for $k = 2, 3, \dots, 17$. One pattern repeated 17 times, and no patterns repeated more than 17 time. The pattern (1,2,1,1; 3,30,39) is the single pattern that repeated 17 times. Due to the 3 ms binning, the delays (3,30,39) actually represent windows, e.g. the first delay of 3 ms represents the interval of 3 to 5 ms, and the second delay represents the interval of 30 to 32 ms.

Table 1: Number of complexity- ≥ 4 patterns which appeared more than once

NUMBER OF REPETITION								
2	3	4	5	6	7	8	9	10
22213	15573	11347	7289	4011	2023	952	387	144
		11	12	13	14	15	16	17
		49	24	8	1	0	0	1

It is intuitive that the most powerful statistics, with respect to an alternative hypothesis supporting the existence of fine-temporal patterns, will come from the tail of the

distribution on numbers of repetitions of patterns of a given complexity. Therefore, we adopted the *maximum number of repetitions of a given complexity* as our statistic, S . For the data at hand, $S(T) = 17$.

200 spike sequences were randomly generated according to the procedure outlined earlier, using an interval size $\Delta = 6$ ms. S was evaluated on each of the Monte Carlo sequences, and the results tabulated in the form of a histogram in Figure 2. The procedure was then repeated, with $\Delta = 15$ ms, and the results superimposed in Figure 2. Based on the maximum number of repetitions observed in the original record (17), we reject the hypothesis of no change in firing rates over 6 ms intervals at a .025 level of significance, and we reject the stronger hypothesis of no change in firing rates over 15 ms intervals at a .005 level of significance.




Figure 2: Histogram of statistics from 200 sampled spike records. The pattern of (1,2,1,1; 3, 30, 39) appeared 17 times in original multi-spike record.

It is perhaps worth noting that the actual times of occurrences of the distinguished pattern (1,2,1,1; 3,30,39) did not correspond, in any obvious way, to the behavior of the monkey.

4 DISCUSSION

Using the method described in this paper, what would it take to obtain an unambiguous answer to the question asked in the introduction, namely what is the temporal resolution at which the nervous system “operates”?

The first step would be to find a statistic $S(T)$ which allowed us to reject the hypothesis that the multi-spike train T has no fine temporal structure. We have presented preliminary results which indicate that statistics constructed from counts of synfire-type events may indeed reveal the existence of structure on the 6-ms time scale. Confirmation of these results on other data sets is required before any conclusions can be drawn.

However, as mentioned in the introduction, the mere existence of structure by no means implies that this structure is actually used by the brain for the purpose of information coding, let alone as a binding mechanism. We are still a long way from such a conclusion and the analysis we conducted so far has not revealed any correlation between the patterns found and behavior.

Of course the general issue of time scales may very much depend on which part of the nervous system is under consideration. It should also be stressed that we have explored only a small fraction of what may be considered reasonable candidate statistics $S(T)$. Further work will depend, to a large extent, on the development of efficient search techniques. These will be necessary if the space of statistics is extended and/or if one wishes to analyze the activity of, say, 100 simultaneously recorded neurons.

Electrophysiology and Behavioral Task

The data used in this study were provided to us by Nicholas G. Hatsopoulos, Catherine L. Ojakangas & John P. Donoghue of the Department of Neuroscience, Brown University (Hatsopoulos et al., 1997). Neuronal activities from the supplementary motor area (SMA) (See Tanji & Shima, 1994, for neuronal activity in SMA) of a monkey performing a delayed sequential task were recorded using a chronically implanted 10×10 -grid multi-electrode array developed at the University of Utah (Nordhausen, 1996).

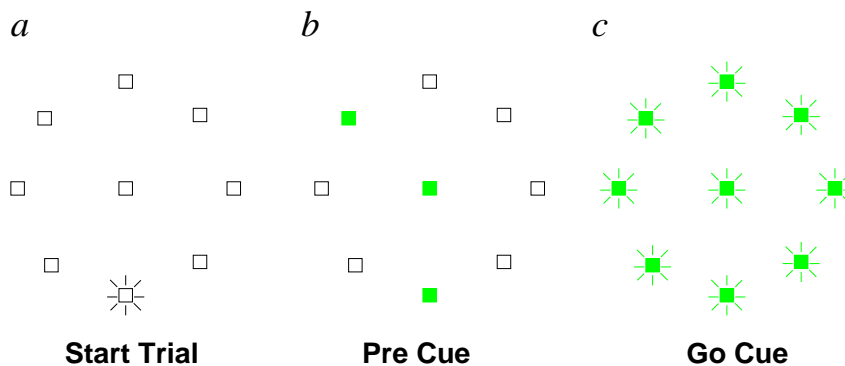


Figure 3: Random spike movements (spike jittering), *a*, original spike record *b,c*, jittered spike records.

During the task, the monkey used a horizontal two-joint manipulandum to move a cursor presented on a video monitor in front of him. Each trial started with the appearance of 9 white squares, with the bottom center square blinking (See Fig.3 *a*). After the monkey held the cursor on the bottom center position (“start position”) for 500 ms, three targets (green) were presented (Fig.3 *b*), two of which were at fixed locations (bottom center and center) and the third at one of two possible locations (upper left or upper right). The monkey was trained to maintain the cursor at the start position until the go cue (all of the square cursors blinking simultaneously) was presented, 1.0 to 1.5 second after the pre-cue. In order to obtain liquid reward, the monkey had to move the cursor, through the center position, to the upper left or right target position.

The data analyzed here used five neurons recorded simultaneously with three electrodes, and lasted about 20 minutes. There were 82 completed trials, of which 79 were

successfully completed and rewarded. The data analyzed included all 82 trials. Out of the 79 successful trials, there were 35 where the monkey had to move the cursor to the upper-right target, and 44 where the monkey had to move the cursor to the upper-left target. The data analyzed lasted from the start of each trial to the end of movement of that trial; thus, all inter-trial periods were ignored.

Acknowledgments

We thank N.G.Hatsopoulos, C.L.Ojakangas & J.P.Donoghue for generously sharing their multi-electrode data. This work was supported in part by National Science Foundation grant SBR-9720368, a Monbusho Grants-in-Aid for Scientific Research, a JSPS Research Fellowships for Young Scientists, Army Research Office contract DAAH04-96-1-0445, National Science Foundation grant DMS-9217655, and Office of Naval Research contract N00014-97-0249.

References

- Abeles, M., Bergman, H., Margalit, E. & Vaadia, E. (1993) Spatio-temporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology* **70**(4):1629-1638.
- Damasio, A.R. (1989) The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation* **1**:123-132.
- Hatsopoulos, N.G., Ojakangas, C.L. & Donoghue, J.P. (1997) Planning of sequential arm movements from simultaneously recorded motor cortical neurons. *Society for Neuroscience Abstract* **23**:1400.
- Livingstone, M.S. (1996) Oscillatory firing and interneuronal correlations in Squirrel monkey striate cortex. *Journal of Neurophysiology* **75**(6):2467-2485.
- Nordhausen, C.T., Maynard, E.M. & Normann, R.A. (1996) Single unit recording capabilities of a 100 microelectrode array *Brain Research* **726**:129-140.
- Tanji, J. & Shima, K. (1994) Role for supplementary motor area cells in planning several movements ahead *Nature* **371**:413-416.
- von der Malsburg, C. (1995) Binding in models of perception and brain function *Current Opinion in Neurobiology* **5**:520-526.