

Rate Limitations of Unitary Event Analysis

A. Roy

P. N. Steinmetz

E. Niebur

Zanvyl Krieger Mind/Brain Institute, Johns Hopkins University, Baltimore, MD 21218, U.S.A.

Unitary event analysis is a new method for detecting episodes of synchronized neural activity (Riehle, Grün, Diesmann, & Aertsen, 1997). It detects time intervals that contain coincident firing at higher rates than would be expected if the neurons fired as independent inhomogeneous Poisson processes; all coincidences in such intervals are called unitary events (UEs). Changes in the frequency of UEs that are correlated with behavioral states may indicate synchronization of neural firing that mediates or represents the behavioral state.

We show that UE analysis is subject to severe limitations due to the underlying discrete statistics of the number of coincident events. These limitations are particularly stringent for low (0–10 spikes/s) firing rates. Under these conditions, the frequency of UEs is a random variable with a large variation relative to its mean. The relative variation decreases with increasing firing rate, and we compute the lowest firing rate, at which the 95% confidence interval around the mean frequency of UEs excludes zero.

This random variation in UE frequency makes interpretation of changes in UEs problematic for neurons with low firing rates. As a typical example, when analyzing 150 trials of an experiment using an averaging window 100 ms wide and a 5 ms coincidence window, firing rates should be greater than 7 spikes per second.

1 Introduction ---

Despite intensive research efforts over half a century, fundamental questions regarding the nature of neuronal representations remain unanswered. One long-standing debate (Hubel, 1959; Werner & Mountcastle, 1963; Smith & Smith, 1965; Griffith & Horn, 1966; Bullock, 1970; Softky & Koch, 1993; Shadlen & Newsome, 1995; Softky, 1995) is whether this representation involves only the mean rate of action potentials, averaged over times on the order of a second, or whether finer temporal structures play a role. Of particular interest are those time structures involving multiple neurons. For instance, synchronous or oscillatory firing of neurons in the cortex has been proposed as a mechanism for representing sensory information (Singer,

1993), for labeling different attributes of a stimulus (Decharms & Merzenich, 1996), and as a possible representation of internal behavioral states (Gerstein & Clark, 1964; Abeles, 1982b, 1991; Dayhoff & Gerstein, 1983a, 1983b; Niebur, Koch, & Rosin, 1993; Niebur & Koch, 1994, Steinmetz et al., 2000). As techniques for recording from multiple neurons become more commonplace, methods for analyzing data from simultaneous recordings become increasingly important. Methods now considered part of the standard repertoire are the cross-correlogram (see Eggermont, 1990, for a review), the joint peristimulus time histogram (Aertsen, Gerstein, Habib, & Palm, 1989), and the gravity method (Gerstein, Bedenbaugh, & Aertsen, 1989). Any addition to the available methods is welcome. There is a particular need for methods that do not require extensive averaging over trials since the brain has to solve problems as they appear, without repetitions and multiple samples of a stimulus. If the brain is able to work without averaging, then it should be possible to detect the "code" that the brain uses by applying algorithmic methods that do not require averaging either. It was one of the goals leading to the development of unitary event analysis to provide such a method (Grün, 1996).

The basic task of unitary event analysis is the detection of periods of synchronous firing of multiple neurons (Grün, Aertsen, Abeles, & Gerstein, 1993; Grün, 1996; Riehle et al., 1997). The method determines intervals of time during an experiment when the number of coincident firings of two or more simultaneously recorded neurons significantly exceeds the number that would be expected if the neurons fired independently. Action potentials occurring simultaneously during these intervals are labeled unitary events (UEs).

Changes in the frequency of UEs correlated with the presumed perceptual state of the animal subject were recently observed in monkey primary motor cortex by Riehle et al. (1997). In these experiments, UEs occurred just prior to the possible appearance of a visual movement cue. More important, the frequency of UEs was significantly higher when the monkey later performed the instructed movement correctly than when the animal failed.

The purpose of this article is to analyze this new method and point out limitations on its use. We do so using theoretical analysis and numerical simulation and by application of UE analysis to data recorded in primate somatosensory cortex. The central result is that any application of UE analysis requires careful examination of the firing rates of the analyzed sequences of action potentials ("spike trains"). While earlier work (Grün, 1996) highlighted limitations of the method for spike trains with high firing rates, we demonstrate here that use of the method for low-rate spike trains (0–10 spikes/s) can also produce artifactual changes in the occurrence of UEs that may appear correlated with behavioral states (as is the case in the data analyzed in this article). This report illustrates this artifact, examines why it occurs, and determines the lower bounds on the firing rates when UE analysis can be used. (Parts of this work have been presented in abstract form in Roy, Steinmetz, & Niebur, 1998.)

2 Unitary Events Method Defined

Unitary events are coincident action potentials that occur at a greater frequency than would be expected if the spike trains were independent (Grün, 1996; Riehle et al., 1997). The basic idea is to count the observed number of coincidences (simultaneous firing of two or more neurons during a small analysis interval of width b , typically 5 ms) in an interval of width T_w (typically 100 ms). This count is compared with the number of coincidences that would occur if each neuron fired independently with the average rate over the interval T_w .

To perform this analysis, each spike train is represented as a binary (0, 1) sequence in time. Each of its segments of length T_w is subdivided into $N_b = \frac{T_w}{b}$ exclusive bins of width b . Each bin has either the value 0, representing no spikes, or 1, representing one or more spikes occurring in the bin. This sequence will be denoted $s(j)$, $j = 1, \dots, N_b$. In UE analysis, it is assumed that spike trains obey Poisson statistics of constant rate within any interval T_w (i.e., that a spike train is a homogeneous Poisson process within this interval). Rates are allowed to be different in different segments, but the firing is always assumed to be governed by Poisson statistics. Therefore, neural firing is assumed to be an inhomogeneous Poisson process over intervals longer than $^1 T_w$.

Under these assumptions, the probability of a neuron's firing at least once in a bin is

$$P(s(j) = 1) \equiv 1 - P(\text{no spikes}) = 1 - e^{-n/N_b}, \quad (2.1)$$

where n is the total number of events (ones) observed in T_w .

Multiple repetitions ("trials") of an experiment are treated by assuming stationarity across the N_t trials but not necessarily within a trial beyond a time T_w . This allows the corresponding intervals from each trial to be concatenated into a single process of length $N_t \times T_w$, which is stationary by assumption. In this case, the number of bins per segment is $N_b = \frac{N_t T_w}{b}$.

¹ Note a conceptual difficulty here that is due to the overlap between neighboring segments. Strictly speaking, the rate in any segment determines the rate in all segments, because of this overlap in conjunction with the requirement of constant rate within any segment. Violation of this assumption may be tolerable as long as the firing rates vary slowly, compared to T_w . In practice, observed firing rates do not always change slowly; for an example, see the raster plots in Figure 1 around the times of stimulus onset and offset, or Figures 2 to 4 in Riehle et al. (1997). Grün (1996) developed more sophisticated methods for the choice of T_w , which may alleviate the practical consequences but do not solve the principal problems addressed later in this article. We will ignore this technical difficulty of UE analysis.

The spiking activity from N simultaneously recorded neurons is described by a joint process composed of N parallel binary processes. The process is represented by a time sequence of N -tuples, $V(j), j = 1, \dots, N_b$. Each component of the N -tuple is either zero (no action potential) or unity (one or more action potentials). Let V_k be a particular N -tuple and let V_{ki} denote its i th component. Assuming that the neurons fire independently, the joint probability of observing the N -tuple V_k is then given by the product of the marginal probabilities for each of the N processes:

$$P(V_k) = \prod_{i=1}^N \left\{ \begin{array}{ll} e^{-n_i/N_b} & \text{if } V_{ki} = 0 \\ 1 - e^{-n_i/N_b} & \text{if } V_{ki} = 1 \end{array} \right\}, \tag{2.2}$$

where n_i is the number of bins with unity value in the concatenated process for neuron i .

Given the assumed stationarity, the number o_k of occurrences of a particular N -tuple V_k in N_b bins has a binomial distribution. Its probability is given by

$$P(o_k|P(V_k), N_b) = \binom{N_b}{o_k} P(V_k)^{o_k} [1 - P(V_k)]^{N_b - o_k}. \tag{2.3}$$

For each number of occurrences o_k of N -tuple V_k , we can compute the significance level for a critical region that contains just this outcome. The significance level will be the total probability of finding the observed number, o_k , of occurrences of V_k or an even larger and less likely number of occurrences. We define this level as the joint p -value of the outcome:²

$$\begin{aligned} \text{joint } p\text{-value}(o_k|P(V_k), N_b) &\doteq \sum_{r=o_k}^{N_b} P(r|P(V_k), N_b) \\ &= \sum_{r=o_k}^{N_b} \binom{N_b}{r} P(V_k)^r [1 - P(V_k)]^{N_b - r}, \end{aligned} \tag{2.4}$$

where $P(V_k)$ is given by equation 2.2.

In the cases considered in this article, only two simultaneous spike trains are analyzed. Therefore only four firing patterns of V_k are possible, and only one of them, $\binom{1}{0}$, representing coincident firing, is of primary interest.

² In the original UE analysis of Grün (1996), the joint p -value is based on the Poisson approximation to the binomial distribution (see her chapter 3). Although this approximation is frequently justified, it fails for the precise considerations introduced in section 3.3, and we therefore use the exact form in equation 2.4.

In the method of UE analysis, if the joint p -value is less than the target p -value, then all coincidences inside the averaging window are labeled unitary events.³ The target p -value (or α level) is the arbitrary significance level chosen to test the null hypothesis that the two simultaneously recorded spike trains are independent. In this article, we consistently use a target p -value of 0.05.

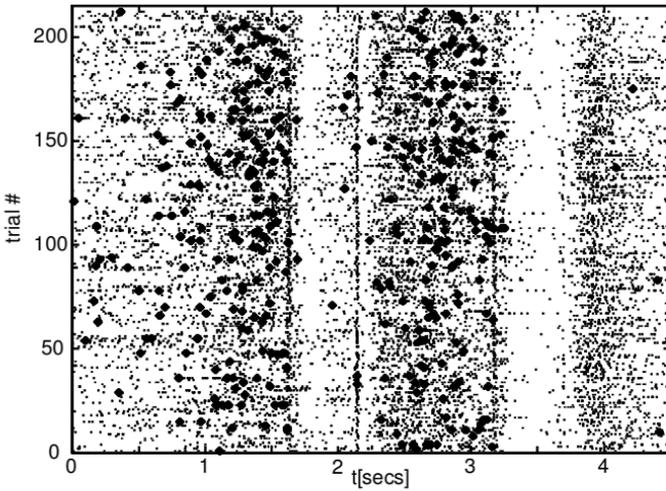
3 Analysis of Unitary Events Method

3.1 Application to Cortical Spike Trains. Unitary events were detected in simultaneously recorded pairs of spike trains originating from individual neurons in somatosensory cortical area II (SII) of macaque monkey. During each trial, a video screen displayed either a blank screen indicating that the monkey was required to perform a tactile task, or a rectangular box consisting of a white outline and black center indicating that the monkey was required to perform a visual task. The tactile task consisted of discriminating the orientation of a 2 mm wide nonresilient bar pushed twice perpendicularly against the distal finger pad of a restrained finger; the orientation of the bar could be the same or different between the two presentations in the same trial. The monkey had to decide whether the orientations were identical or different. The visual task consisted of distinguishing whether the left or right box presented on the video screen was dimmed. In both cases, the monkey indicated its response by pressing a foot switch. Details of the experimental procedure were as described by Hsiao, O'Shaughnessy, and Johnson (1993) except that the tactile and visual stimulation were as described above.

Figures 1A and 1B show recorded action potentials of two neurons in SII as well as UEs during 212 trials of this experiment while the monkey

³ A possible alternative, suggested by Grün (1996), is to label as unitary events only some of the coincidences in the interval. Their number would be given by the number of coincidences *exceeding* that expected to occur by chance given the observed firing rates. This is not the case in the standard UE analysis approach as defined by Grün (1996) and Riehle et al. (1997) in which all coincidences are labeled UEs independent of the firing rates. We cannot identify which ones are in excess of the total number of coincidences. We can then make a statement that $x\%$ of them are in excess of what is expected. Alternatively, the coincidences that qualify as UEs could be selected randomly at the appropriate frequency. Another approach could be to generate a weighting procedure for the UEs (i.e., the weight of a coincidence increases monotonically (linearly) if it is detected as being significant in multiple overlapping blocks). After the completion of the fixed number of overlapping tests for each UE, the ones with the least weight are removed until the theoretical expected number is reached.

A



B

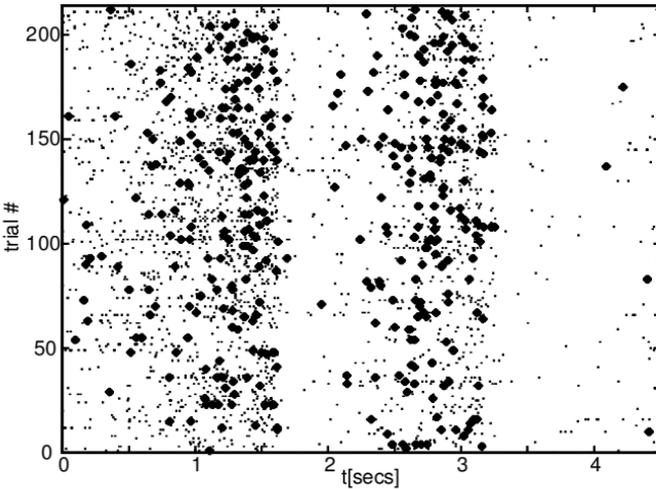


Figure 1: (A,B) Raster plots of action potentials (dots) and UEs (diamonds) in two cells as a function of time within each trial when the monkey performed the tactile task, requiring attention to the task. Each row corresponds to one trial. The stimulus bar is in contact with the skin from about $t = 1.7$ s to 2.2 s and from $t = 3.3$ s to 3.8 s; the “on” and “off” responses can be clearly seen for the cell shown in A. See the text for details.

worked on the tactile task. Note that the UEs are identical in parts A and B of Figure 1 since they are defined by the coincident firing of the two neurons whose action potentials are shown in these two figures. Note also that the assumption of stationarity across (but not within) trials appears justified. Figures 1A and 2A contrast the data that were obtained when the monkey worked in the tactile and visual task, respectively, in the same neuron. The intervals between 1–1.5, 2.5–3.0, and 3.8–4.3 seconds show a change in the frequency of UEs when the intervals are compared between Figures 1A and 2A, thus demonstrating that the frequency of UEs is correlated with the behavioral state of the primate. The stimulus bar contacts the skin during the time intervals between $t = 1.7$ s and 2.2 s and between $t = 3.3$ s and 3.8 s. Note that the UE frequency increases before (and not during) the stimulation. The underlying joint p -values are shown in Figure 2D, which plots (for better resolution) $\log_{10} \frac{1-p}{p}$ for the tactile task. A simple explanation assuming that UEs are suppressed by the application of the stimulus bars does not hold, since UE frequency decreases at the end of the trial, when the stimulus bar is again removed from the skin (after $t = 3.8$ s). Note, however, that the firing frequency of neuron 2 (raster plots in Figure 1B) is reduced toward the end of the trial (after about 3.2 s). Although UE analysis is designed to compensate for rate effects, we will see that changes in the firing rate may lead to changes in UE frequency.

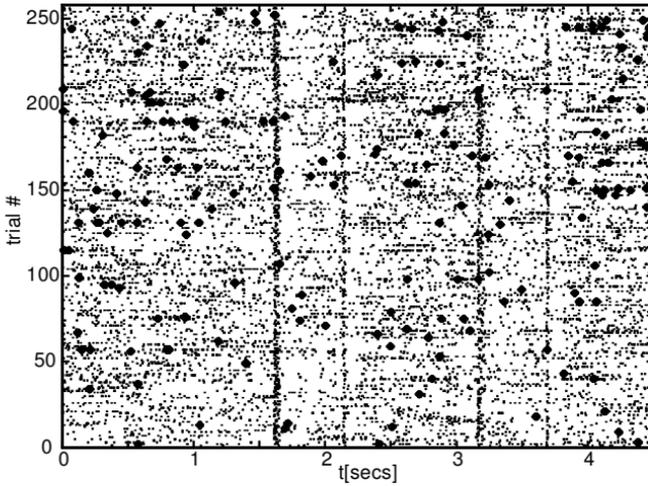
Figure 2A shows that no such increase in UE frequency is observed when the animal is performing the visual task in the same intervals with identical sensory stimulation (remember that the same tactile stimulus is being delivered to the monkey even when he is doing the visual task). Increased UE frequency therefore could be interpreted as signaling a buildup of an anticipatory state in the animal rather than a response to sensory stimulation.

Of the 30 pairs of neurons examined, 9 (30%) showed changes in UE frequency correlated with the behavioral state of the monkey. Recordings from an additional 7 cell pairs contained many UEs, but UE occurrence did not vary with behavior. The remaining 16 pairs contained an insignificant number of UEs.

These results appear comparable to those reported by Riehle et al. (1997) in motor cortex, where the frequency of UEs increases sharply in those intervals in which the monkey might anticipate the occurrence of a visual cue for movement. We will see that this interpretation must be made cautiously and may not be warranted, particularly not with low firing rates.

3.2 Discrete Distribution of the Number of Coincidences and Possible Significance Levels. In order to elucidate the mechanisms underlying the changes of UE frequency, we reevaluated the statistical test used to identify UEs. As reviewed in section 2, the method of UE analysis detects coincident spikes occurring in time intervals where the null hypothesis of independent neural firing is rejected. This interpretation of changes in UE frequency will

A



B

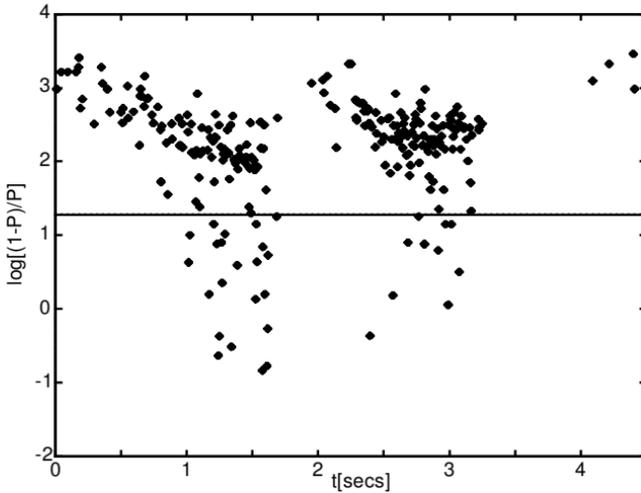


Figure 2: (A) Similar raster plot as in Figure 1A when the monkey performed the visual task, which does not require tactile attention. (B) Joint p -values (P in the graph) plotted as a ratio of logarithmic functions for overlapping 100 ms time segments in the tactile task. The horizontal line at ordinate value $\log_{10} 19 \approx 1.3$ shows the target significance level (in this case, 0.05). At times (i.e., values on the abscissa) without entries in the figure, bins have no recorded spikes in one or both neurons.

be valid, however, only if the significance level of the test for independent firing is equal for the time intervals being compared. If the significance level of the test varies, then changes in UE frequency are due not only to changing frequencies of coincidences but also to changing significance levels. It is a statistically ill-defined procedure to compare UE frequencies between intervals with different significance levels.

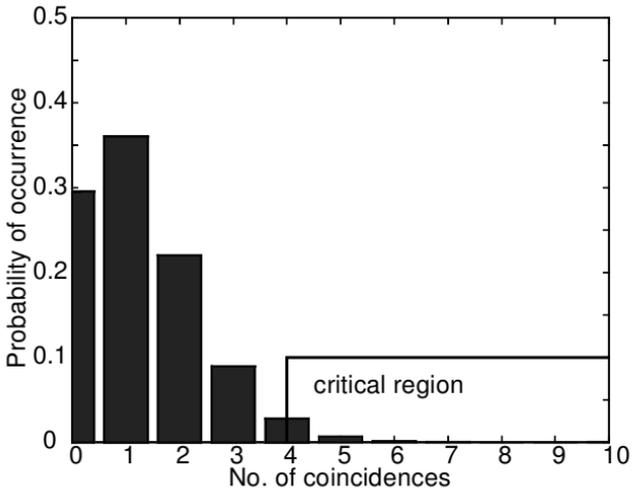
The significance level of the test of independent firing is equal to the probability of a time interval containing a number of coincidences that lies in the critical region of the test. The structure of this critical region is determined by the probability distribution of the number of coincidences for independent processes. As reviewed in section 2, this distribution is binomial, with its shape depending on both the number of time bins examined and the rates of neural firing. Two examples with 2000 bins and two different rates are shown in Figures 3A and 3B.

Critical regions for the maximum likelihood test used in UE analysis consist of right-hand tails of this distribution, since the maximum likelihood rule requires that if a specific number of coincidences is included in the critical region, then all less likely coincidences must also be included (Lindgren, 1976). It is therefore important to consider the effective significance level, which we define as the total probability of the largest critical region that does not exceed the target significance level p . It is usually obtained by integrating the right tail of the probability density function leftward from large numbers until the area reaches the target p value.

Examples are shown in Figure 3A for a firing rate of 5 spikes per second and in Figure 3B for 20 spikes per second. The figures show the critical regions closest to, but not exceeding, a significance level of $p = 0.05$. The crucial observation is that the effective significance levels are discrete, since the number of coincidences in a window of given size T_w is an integer and therefore a discrete variable. At low firing rates, such as 5 spikes per second, this discrete distribution creates effective significance levels that are broadly spaced and usually approximate the target level poorly. The effect is less important for high firing rates, such as 20 spikes per second (for the chosen parameters), where a significance level near the target significance level can be chosen with reasonable accuracy (see below for a more quantitative statement). As the mean of the binomial distribution increases (both with increasing bin size and with increasing average firing rate), the effective significance level approaches p , but for many physiologically realistic spike rates, the difference between the effective significance level and p itself is a substantial fraction of p .

The dependence of the effective significance level on neuronal firing rate is shown in Figure 4, which plots the effective significance level for a pair of independent neurons using a target $p = 0.05$ for 3000 time bins of length 5 ms each (this corresponds to a typical experiment with $N_T = 150$ trials and $T_w = 100$ ms). The most problematic point when performing UE analysis is the presence of large jumps in significance level as a function of rate. These

A



B

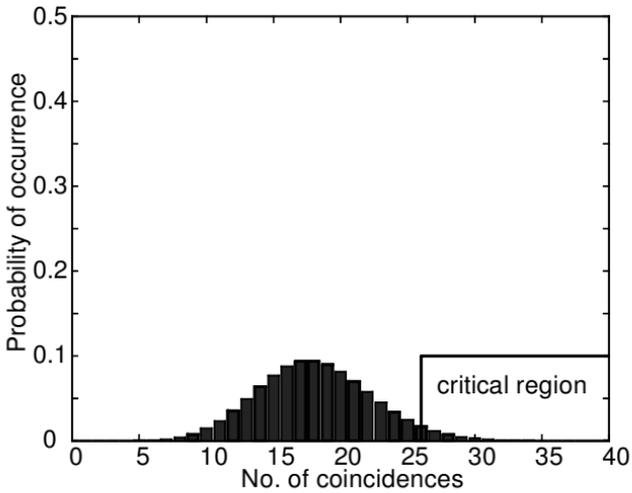


Figure 3: (A) Probability distribution of the number of coincidences for a firing rate of 5 spikes per second. The effective significance is shown for a significance level of $p = 0.05$. (B) As in (A) but for 20 spikes per second.

occur when a particular outcome, such as three coincidences per bin, is no longer included in the critical region as the rate increased. It can be seen that a small change in firing rate (e.g., 0.1 spike/s) can lead to a very large change in the size of the critical area—the effective significance level. The

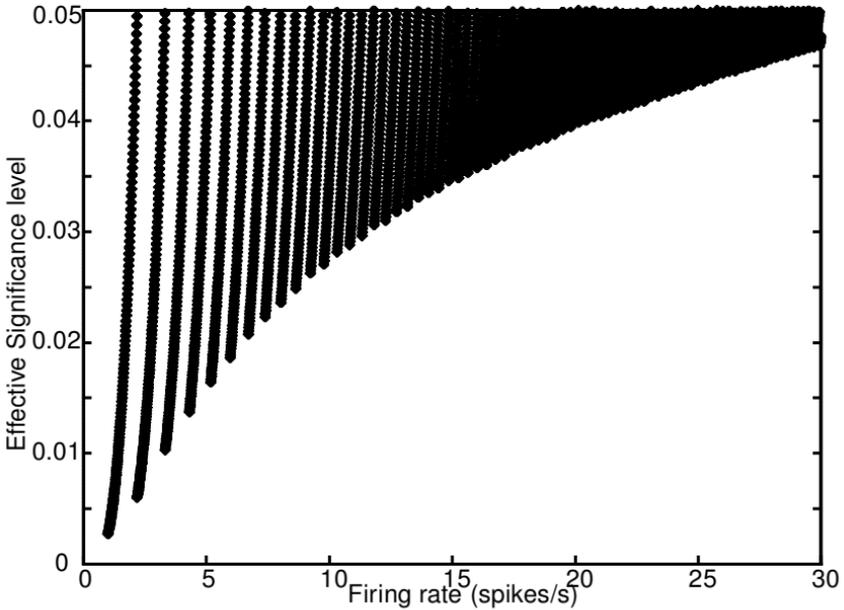


Figure 4: Effective significance levels for the maximum likelihood test of independent firing. $N_b = 3000$, $b = 5$ ms.

largest jumps appear at small firing rates. For instance, in the case illustrated in Figure 4, a small increase in firing rate at about 2 spikes per second (first branch shown in the figure) leads to an increase in effective significance by a whole order of magnitude, from $p = 0.05$ to about $p = 0.005$. It should be noted that the firing rate of the model can change by an arbitrarily small amount, although the smallest change in the observed rate in a simulation will correspond to the addition or removal of a single spike over N_T trials in the averaging window of width T_w .

We have thus determined that the discrete nature of coincident events is the cause of the observed large changes in significance level. Small changes of neural firing rates correlated with behavioral status may therefore lead to large changes in significance level. If this is the case, we cannot conclude that different frequencies of UEs are due to differences in the correlation state of the neurons. In terms of UE analysis, which is based on constant significance levels, the apparent changes in UE frequency might represent artifacts of the method.

Although the jumps in effective significance level for a single test are large, we show in the next section that they do not appear clearly in UE analysis because the method includes testing the coincidences multiple times. The effects of this repeated testing will be explored next.

3.3 Rate Dependence of Unitary Events. As described in section 2, the presence of UEs is determined by the number of coincidences within a window of length b in an interval T_w . If this number is in the critical region of a test for independent firing, then all coincidences within the interval are considered UEs. After each test, the interval slides by the width of the coincidence window, b , and the test is repeated. Therefore, each coincidence is subjected to T_w/b dependent tests, which means 20 tests for typical window sizes $T_w = 100$ ms, $b = 5$ ms. Furthermore, the average rate estimate used to construct a single test is a random variable. Since the sliding window moves one bin at a time, the tests of the null hypothesis are dependent.

The randomness of the rate estimate and the interdependence of these tests make it difficult to determine analytically the overall significance level of the test for coincident firing employed in UE analysis. We therefore compute the significance level by Monte Carlo simulation. The significance level is equal to the probability that coincident spikes are labeled as unitary events even though the null hypothesis of independent firing is true. The average rate estimate for a single test was simulated as being equal to $(T_w/b)^{-1} \sum_{\alpha=1}^{T_w/b} Y_\alpha$, where Y_α are independent Poisson-distributed random variables drawn from a fixed firing rate within the time window of width T_w .

The Monte Carlo computations consisted of 10,000 simulated trials of $3N_b - 1$ bins each for two independent Poisson processes at each value of the firing rate. Unitary events were detected using the methods described in section 2. The frequency of UE occurrence was defined as the number of computed UEs in the central block of N_b columns of bins divided by bN_bN_t (this avoids edge effects). The mean and variance of the UE frequency were then computed over all iterations of the simulation and shown in Figure 5.

This graph shows that the average frequency of UE occurrence is a monotonically increasing function of rate. Intuitively, this can be understood from the fact that the total number of accidental coincidences increases with the product of the firing rates of the two neurons and that therefore the number of accidental coincidences labeled as UEs also increases.

The dependence of UE frequency on firing rates has previously been noted by Grün (1996) at higher firing rates, and it can lead to marked increases in UE frequency. For example, in Figure 5, it can be seen that a twofold rate change from 4.5 spikes per second to 9 spikes per second causes more than a fourfold increase in UE frequency.

In addition, it is important to notice that 20 dependent tests are performed for every coincidence since each bin of width $b = 5$ ms is member of 20 windows of width $T_w = 100$ ms, each of which has, in general, a different firing rate. If any of these tests detects a significantly elevated number of coincidences, all coincidences in the interval are labeled UEs. This causes the transformation of the "sawtooth" function in Figure 4 into the smooth increase in average number of UEs shown in Figure 5. The large jumps in

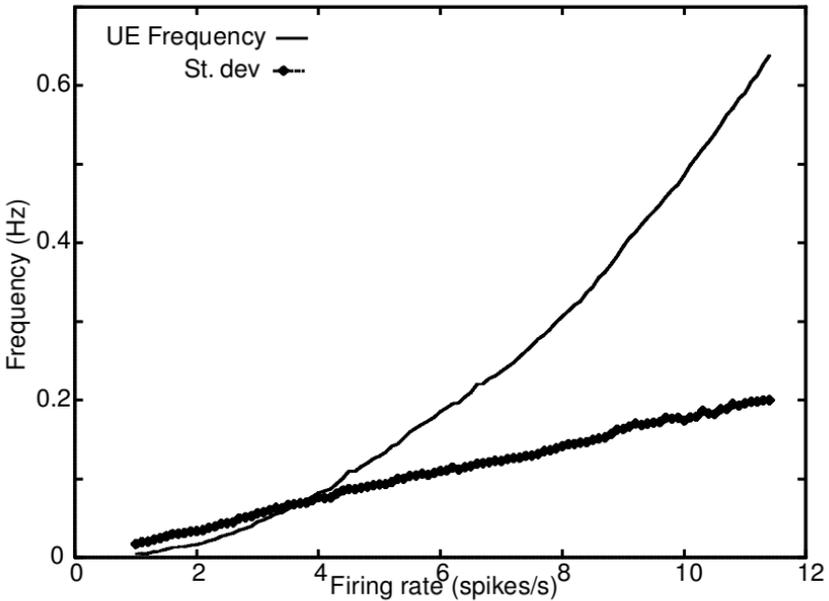


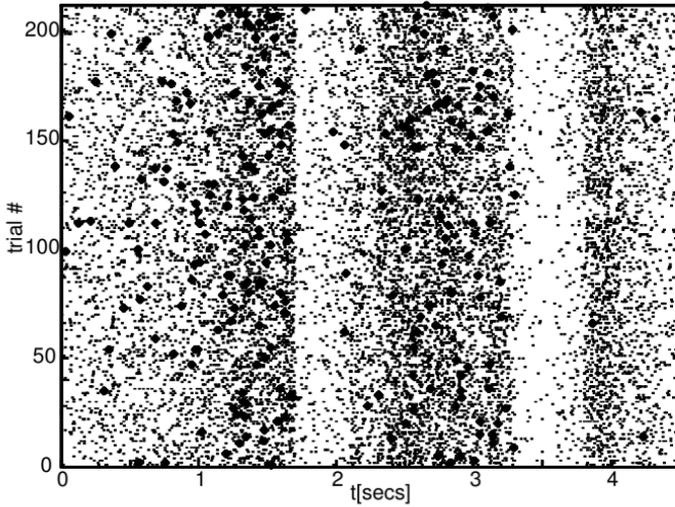
Figure 5: UE frequency as a function of firing rate. $N_b = 3000$, $b = 5$ ms.

significance levels in Figure 4 are reflected in the large standard deviation of UE frequency relative to its mean shown in Figure 5. Since the size of the jumps decreases with increasing firing rates, the same follows for the standard deviation of UE frequency relative to its mean.

Thus, although the target significance level used in performing UE analysis remains constant, the effective significance level of the test for unitary events is dependent on firing rate. In section 3.4 we will develop criteria for deciding when the frequency of UEs is a reliable quantity. Before doing this, however, we applied this analysis to the data recorded in somatosensory cortex.

Figure 6 demonstrates the effect of varying firing rates on UE frequency. We computed the time-varying firing rates for each of the two neurons whose activity is plotted in Figures 1A and 1B using a 100 msec averaging window and generated two (nonhomogeneous) Poisson processes using these rate profiles. The time-varying rate profiles were computed consistent with the assumed null hypothesis for UE analysis (which is in this case homogeneous or stationary rates over 100 msec). UE analysis was then applied to the simulated processes. Figure 6A shows the UEs obtained. The pattern of UEs is remarkably similar to that observed in Figure 1A. Similarly, the logarithmic plot of the significance of the events shown in Figure 6B captures the relevant features of the corresponding plot in Figure 2B.

A



B

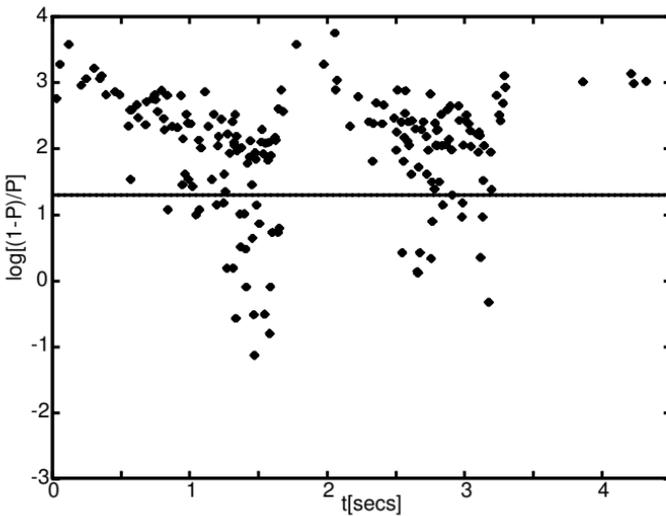


Figure 6: (A) Raster plots of action potentials (dots) and UEs (diamonds) in two cells simulated as realizations of two independent inhomogeneous Poisson processes with a rate profile similar to the cell from Figure 1A. See the text for details. (B) Joint p -values (P in the graph) plotted as a ratio of logarithmic functions for overlapping 100 ms time segments.

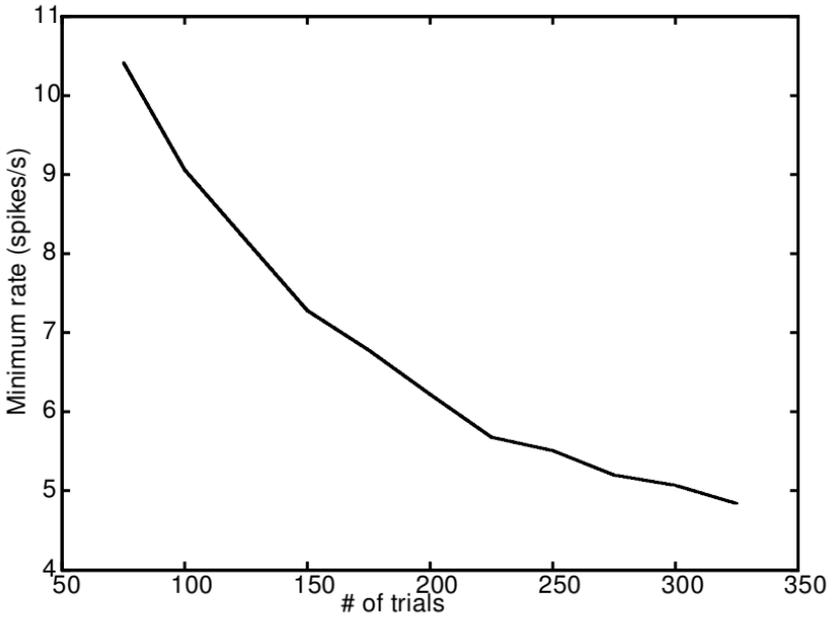


Figure 7: Minimum rates where mean UE frequency = 2σ .

A uniformly distributed (or random) pattern of UEs was obtained when we applied UE analysis to similarly generated inhomogeneous spike train simulations of each of the two neurons whose UE analyses are shown in Figure 2A. The random nature of the distribution of UEs matches the obvious lack of pattern in Figure 2A (data not shown). Therefore, based on these simulations, it can be argued that a large portion of the observed pattern of UEs in those figures is due to rate effects and not to rate-independent coincidences.

3.4 Firing Rates with Reliable Significance Levels. Figure 5 shows that the frequency of UE occurrence for low neural firing rates is a random variable with a large standard deviation relative to its mean under the null hypothesis of independent firing. This random variation of the mean UE frequency makes it difficult to interpret changes in UE frequency. If this variation is high relative to the mean significance level, then changes in UE occurrence may be due to changes unrelated to changes in the correlational state of the neurons. For instance, changes in the firing rate that are orders of magnitude smaller than the standard deviation of the average firing rate (fractions of a Hz; see Figure 4) can change the effective significance level by a factor of 10. As a lower bound on rates where the variation is acceptable, we show in Figure 7 the minimum rate where a

95% confidence interval excludes zero or, equivalently, where the mean is equal to twice the standard deviation. Other intervals may be appropriate depending on the degree of confidence desired in a specific investigation. For example, when analyzing simulated data for a typical experiment comprising 3000 time bins (100 ms averaging and 5 ms analysis windows for 150 trials), a 95% confidence interval (2σ) for the significance level includes zero for firing rates below 7.3 spikes per second. As the number of trials increases, lower rates are sufficient to produce a reliable significance level.

4 Discussion

It is well established that the average rate of neuronal firings is used to represent information in parts of the nervous system (Kandel, Schwartz, & Jessel, 1991). It is likely that the nervous system also makes use of some of the temporal variation of the firing rate, for example, of onset and offset responses of some neural populations. It is a matter of active discussion (and has been for several decades; Hubel, 1959; Werner & Mountcastle, 1963; Smith & Smith, 1965; Griffith & Horn, 1966; Bullock, 1970) whether more subtle forms of temporal structures are being employed to encode and process neuronal information.

One of the attractive features of temporal codes is that they do not rely on averaging over either large populations or long time intervals. Instead, it is conceivable that very specific patterns are generated in one part of the nervous system and instantaneously decoded and processed in other parts. As long as it is highly unlikely that such patterns are generated by chance, little or no averaging is required to recognize the patterns with a high degree of confidence. Codes that operate without extensive temporal averaging have clear advantages: the brain of a behaving animal needs to solve problems as they appear, without waiting for repetitions and multiple samples of a stimulus.

If this is the case, it should be possible to detect such temporal structures by algorithmic methods that do not require extensive averaging over either time or population. This is not the case for many of the methods used to study the temporal structure of neuronal responses. For instance, the autocorrelation function and its Fourier transform, the power spectrum, are global (nonlocal) measures in the temporal domain. The same is true for their generalizations to multiple spike trains, that is, the cross-correlation function and its Fourier transform. Therefore, such methods require data on the firing behavior over a long time period and are less suitable for making quick decisions.

A recently proposed local method that does not rely on averaging over long time periods is unitary event analysis (Grün et al., 1993; Grün, 1996; Riehle et al., 1997). The underlying assumption is the existence of UEs in brain activity, represented by "brief unusual constellations of activity"

among neuronal assemblies (Grün, 1996). It is assumed that the degree of unusuality is given by the degree of interneuronal correlation and measured by a quantity closely related to the surprise measure used previously in neurophysiology (Lélegendy, 1975; Lélegendy & Salcman, 1985; Palm, Aertsen, & Gerstein, 1988) whose expectation value is, in turn, closely related to Shannon's (1948) information and the entropy of statistical physics (Boltzmann, 1887). Since the frequency of accidental correlations increases with the firing rates of participating neurons, applying appropriate corrections requires knowledge of mean firing rates, which are computed by averaging over trials.

Grün (1996) has pointed out that UE analysis is of limited value if the firing rates of the analyzed neurons are high. The reason is, essentially, that the number of accidental coincidences becomes too large, which makes the detection of possible nonaccidental coincidences more difficult. In this article, we report limitations of this method that are particularly stringent in the case of *low* neural firing rates. This is of particular relevance since the firing rates in the more central cortical areas are usually low and similar to those found to be problematic in our analysis (a few impulses per second or less; see, e.g., Abeles, Vaadia, & Bergman, 1990). Since it is usually assumed that temporal codes such as those assumed in UE analysis are of more relevance in central than in peripheral cortical areas, the limitations reported here are pertinent.

It is also important to note that UE analysis does not provide a measure of significance of the absolute number of excess coincidences. It is therefore unlike the significance of spatiotemporal patterns computed by Abeles (1982a) and Abeles, Bergman, Margalit, and Vaadia (1993). Rather, UE analysis attaches a label to brief periods of increased coincident firing and examines changes in the pattern of the label that are correlated with different behavioral or response states. In fact, the overlapping nature of the test epochs used in UE analysis makes it difficult to determine the significance of the number of excess coincidences as a whole.

At the root of the problem are the underlying discrete statistics and the resulting changes in effective significance levels of the test for coincident firing. We show that for low but realistic firing rates (e.g., 1 Hz), the effective significance level for coincidences labeled as UEs varies considerably. For instance, the significance level of coincidences can be as low as 0.005 instead of the target level 0.05 (see Figure 4)—a whole order of magnitude below the target level—and the average significance level of UEs at low rates is about half of that of the target level. These changes cause a large error in estimation of the mean number of UEs expected under the null hypothesis. Thus, it is difficult to determine whether changes in observed UE frequency are due to random variations or whether they signify changes in actual neural synchrony.

In conclusion, we have pointed out limitations in the applicability of UE analysis in the case of low firing rates. These limitations arise due to

the discrete character of the underlying point process model. If experimental conditions permit the collection of sufficiently large amounts of data, these limitations can be overcome. We provide practical guidelines for the parameter ranges within which the analysis is applicable.

Acknowledgments

We thank Paul Fitzgerald, Steve Hsiao, and Ken Johnson for providing the data presented in Figures 1 and 2. We also thank both anonymous referees for valuable suggestions. This work was supported by the Lucile P. Markey Foundation, the National Science Foundation, and the Alfred P. Sloan Foundation through a Sloan Fellowship to E. N.

References

- Abeles, M. (1982a). *Local cortical circuits*. New York: Springer-Verlag.
- Abeles, M. (1982b). Role of the cortical neuron: Integrator or coincidence detector? *Israel J. Med. Sci.*, *18*, 83–92.
- Abeles, M. (1991). *Corticonics—Neural circuits of the cerebral cortex*. New York: Cambridge University Press.
- Abeles, M., Bergman, H., Margalit, E., & Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiology*, *70*(4), 1629–1638.
- Abeles, M., Vaadia, E., & Bergman, H. (1990). Firing patterns of single units in the prefrontal cortex and neural network models. *Network: Computation in Neural Systems*, *1*, 13–25.
- Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K., & Palm, G. (1989). Dynamics of neuronal firing correlation: Modulation of effective connectivity. *J. Neurophysiol.*, *61*(5), 900–917.
- Boltzmann, L. (1887). Über die mechanischen Analogieen des Zweiten Hauptsatzes der Thermodynamik. *J. Reine Angew Mathematik*, *100*, 201–212.
- Bullock, T. H. (1970). The reliability of neurons. *J. Gen. Physiol.*, *55*, 565–584.
- Dayhoff, J., & Gerstein, G. (1983a). Favored patterns in spike trains. I. Detection. *J. Neurophysiol.*, *49*, 1334–1348.
- Dayhoff, J., & Gerstein, G. (1983b). Favored patterns in spike trains. II. Application. *J. Neurophysiol.*, *49*, 1349–1363.
- Decharms, R. C., & Merzenich, M. M. (1996). Primary cortical representation of sounds by the coordination of action potential timing. *Nature*, *381*, 610–613.
- Eggermont, J. J. (1990). *The correlative brain: Theory and experiment in neural interaction*. New York: Springer-Verlag.
- Gerstein, G., Bedenbaugh, P., & Aertsen, A. (1989). Neuronal assemblies. *IEEE Transactions on Biomedical Engineering*, *36*(1), 4–14.

- Gerstein, G., & Clark, W. (1964). Simultaneous studies of firing patterns in several neurons. *Science*, *143*, 1325–1327.
- Griffith, J. S., & Horn, G. (1966). An analysis of spontaneous impulse activity of units in the striate cortex of unrestrained cats. *J. Physiology (London)*, *186*, 516–534.
- Grün, S. (1996). *Unitary joint-events in multiple-neuron spiking activity*. Frankfurt am Main: Verlag Harri Deutsch.
- Grün, S., Aertsen, A., Abeles, M., & Gerstein, G. (1993). Unitary events in multi-neuron activity: Exploring the language of cortical cell assemblies. In N. Elsner & M. Heisenberg (Eds.), *Gene-brain-behavior* (p. 470). New York: Georg Thieme Verlag.
- Hsiao, S. S., O'Shaughnessy, D. M., & Johnson, K. O. (1993). Effects of selective attention on spatial form processing in monkey primary and secondary somatosensory cortex. *J. Neurophysiology*, *70*(1), 444–447.
- Hubel, D. H. (1959). Single unit activity in striate cortex of unrestrained cat. *J. Physiology (London)*, *147*, 226–238.
- Kandel, E., Schwartz, J., & Jessel, T. M. (1991). *Principles of neural science* (3rd ed.). New York: Elsevier.
- Légyndy, C. (1975). Three principles of brain function and structure. *International Journal of Neuroscience*, *6*, 237–254.
- Légyndy, C., & Salzman, M. (1985). Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *J. Neurophysiol.*, *53*, 926–939.
- Lindgren, B. W. (1976). *Statistical Theory* (3rd ed.). New York: Macmillan.
- Niebur, E., & Koch, C. (1994). A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. *Journal of Computational Neuroscience*, *1*(1), 141–158.
- Niebur, E., Koch, C., & Rosin, C. (1993). An oscillation-based model for the neural basis of attention. *Vision Research*, *33*, 2789–2802.
- Palm, G., Aertsen, A., & Gerstein, G. (1988). On the significance of correlations among neuronal spike trains. *Biol. Cybern.*, *59*, 1–11.
- Riehle, A., Grün, S., Diesmann, M., & Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, *278*, 1950–1953.
- Roy, A., Steinmetz, P., & Niebur, E. (1998). Rate dependence of unitary event analysis. *Soc. Neurosci. Abstr.*, *28*, 1513.
- Shadlen, M. N., & Newsome, W. T. (1995). Is there a signal in the noise? *Current Opinion in Neurobiology*, *5*, 248–250.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.*, *27*, 379–423, 623–656.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annu. Rev. Physiol.*, *55*, 349–374.
- Smith, D. R., & Smith, G. K. (1965). A statistical analysis of the continuous activity of single cortical neurones in the cast anaesthetized isolated forebrain. *Biophys. J.*, *5*, 47–74.
- Softky, W. (1995). Simple codes versus efficient codes. *Current Opinion in Neurobiology*, *5*, 239–247.

- Softky, W., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, *13*(1), 334–350.
- Steinmetz, P. N., Roy, A., Fitzgerald, P. J., Hsiao, S. S., Johnson, K. O., & Niebur, E. (2000). Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*, *404*, 187–190.
- Werner, G., & Mountcastle, V. B. (1963). The variability of central neural activity in a sensory system and its implications for the central reflection of sensory events. *J. Neurophysiol.*, *26*, 958–977.

Received December 4, 1998; accepted September 17, 1999.