Journal of Neuroscience Methods, 45 (1992) 1–14

NSM 01401

An efficient method for detecting connectivity in neural ensembles

Brent W. Edwards [1] and Gregory H. Wakefield

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122 (USA)

(Received 24 December 1991)
(Revised version received 15 June 1992)
(Accepted 19 June 1992)

Key words: Computer analysis; Connectivity detection; Interval statistics; Models; Neural ensemble; Signal processing

Modern technology is allowing researchers to collect data from neural ensembles with a large number of units, and the analysis of interaction between these units can be very time consuming. Estimation of pairwise connectivity is the most common method of determining the neural 'network' but usually necessitates the production of numerous histograms for each pair considered. We present a method which will indicate which pairs in a network represent potential connections and thereby simplify the postexperimental analysis. The technique uses cross-interval information to create an $n \times n$ matrix which represents all possible connections in an n neuron ensemble and can be calculated recursively on-line. The performance of this technique is analyzed with respect to data size and strength of the connections. It is compared to 2 similar techniques that are also presented here, one in which perfect knowledge of the timing of the excitation is known, and one in which the timing can be bounded.

Introduction

The advent of advanced multi-unit recording capabilities has emphasized techniques used to analyze neural ensembles. The study of an ensemble's response to stimuli, rather than an individual neuron's response, is the next step towards understanding how the brain recognizes and classifies information. If a group of neurons encodes information as an ensemble, rather than as a collection of individual units, they must interact with each other and alter each other's firing pattern. Thus, given the firings from a multi-unit recording, methods are needed to analyze the nature of the neural interactions and to determine the connections of the neural ensemble. Not only will more information be revealed about the processing of stimuli but also about the dynamic working of the brain.

Most of the non-parametric techniques developed for analysis of neural ensembles are for pairwise interactions between units within the ensemble (Aertsen and Gerstein, 1985; Melssen and Epping, 1987; Palm et al., 1988; Aertsen et al., 1989; Voigt and Young, 1990) rather than n-wise interactions. This is primarily due to the technical limitations of recording from more than 1 unit simultaneously. Another reason is the difficulty of displaying information about multiple units on a 2-dimensional plot. Cross-correlation functions (Perkel et al., 1967) and scatter plots (Gerstein and Perkel, 1972) are easily displayed and simple to analyze visually when dealing with 2 neurons. Perkel et al. (1975) developed a method for displaying a scatter plot of 3 neurons

Correspondence and *present address:* Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA. Tel.: (612) 625-8557; FAX: (612) 626-2079.

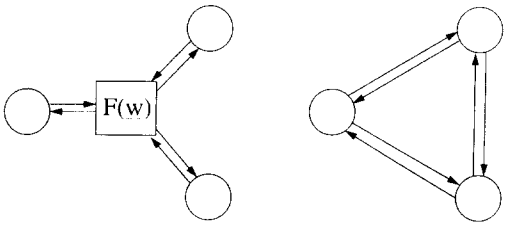

(a) (b)

Fig. 1. a: generalization of neural connectivity, allowing for high-order interaction. b: simplification of neural connectivity where only pairwise interaction is allowed.

in 2 dimensions using what was termed a snowflake plot but, in general, excitatory and inhibitory interactions between units are difficult to represent visually for groups of 3 or more neurons.

If the occurrence times of neuron i are defined by the vector w_i, then the firing statistics of an ensemble of n neurons is fully given by $p(\{w_i\}_1^n)$. This reduces to the more manageable $\prod_{i=1}^n p(w_i)$ only when the neurons are firing independently. In the absence of such independence, the probability distribution for neural firings given neural interaction is difficult to estimate empirically because of the large amount of data needed to define the probability space $p(\{w_i\}_1^n)$ with reasonable confidence. In an effort to reduce complexity, ensemble interactions have thus been reduced to pairwise interactions, reducing the probability space to $p(w_i,w_j|\{w_k\}_{k\neq i,j})$, $i \neq j$. This reduction for the case of 3 neurons is shown in Fig. 1, where (a) depicts the most general model allowing for high-order interaction, defined by F(w), and (b) shows the reduction to pairwise interaction.

This reduction in connectivity order has a significant computational effect on the analysis of large ensembles. If there are n units that are recorded, then $n(n-1)/2$ pairs must be analyzed to examine all possible interactions. If the analysis of each pair requires more than a very simple and quick analysis, the process of detecting and estimating neural interactions for the whole ensemble is very time consuming and increases approximately proportionally to n^2. Parametric techniques exist which have been developed for neural ensembles analysis of which is not reduced to pairwise analysis (Gerstein et al., 1985; Chornoboy et al., 1988; Brillinger, 1988), but the methods for estimating the interactions are generally complex and require much more computational effort than non-parametric methods.

With recent advances in technology allowing data collection for large numbers of units (BeMent et al., 1986; Blum et al., 1991), new techniques are needed to reduce the $n(n-1)/2$ complexity problem. A technique which detects neural interactions but does not necessarily estimate the effect would be advantageous, if it were computationally simple. Such a technique would indicate which pairs should be analyzed so that the number of pairwise relationships considered, and thus computations and time, would be considerably reduced.

In this paper, a method which detects neural interactions will be developed. It is a computationally simple technique which assigns a scalar to each possible connection, and if that number is within a certain decision region then a connection is detected. The information can be represented in the form of an $n \times n$ matrix which is updated with each spike occurrence. The entries are updated recursively and involve few calculations, so this technique can be implemented online, concurrent with the experiment. The performance of the detector will be analyzed with respect to the strength of the connection and the data sized used.

Background theory

An intensity theory of point processes will be used to model and study neuronal behavior. If w_t is a vector which represents the firing times of a neuron up to time t, then the *instantaneous intensity* is defined by

$$\mu(t; w_t) = \lim_{\Delta t \to 0} \frac{\Pr\{\text{firing in } [t,t+\Delta t) \mid w_t\}}{\Delta t}. \tag{1}$$

This represents the firing rate of a neuron at time t conditioned upon a specific realization defined by w_t. For example, if a neuron has a refractory period then $\mu(t; w_t)$ will change according to the firing times w_t.

The *average intensity* is defined by

$$\lambda(t) = \mathrm{E}[\mu(t; w_t)] \tag{2}$$

and represents the unconditional firing rate, i.e., the firing rate of a neuron at time t without knowledge of the previous firing times. This intensity model using μ and λ assumes that the statistics of the firing times are equivalent to that of an inhomogeneous Poisson process, and given the intensity all of the properties of the Poisson process follow. For example, the expected probability of the neuron firing within a small interval Δ about time t is approximately $\lambda(t)\Delta$ for $\lambda(t)\Delta < 0.1$ (Edwards and Wakefield, 1990). See Snyder and Miller (1991) for a detailed description of other properties.

Two properties of point processes will be of concern in this paper. One is the *interval function*, $I(t)$. This is the probability distribution function (pdf) for the time between 2 successive firings. For a homogeneous process, one with a constant $\mu(\cdot)$, the interval function is

$$I(t) = \mu e^{-\mu t}. \tag{3}$$

Throughout the rest of the paper, this will be termed the self-interval function to distinguish it from the cross-interval function.

The cross-interval function, $CI_{AB}(t)$, is the pdf for the time interval between a spike from neuron A and the next spike in neuron B. It can be shown that the pdf for the duration from a random point in time to the next firing of a process is

$$\frac{1 - F(t)}{\mu}, \tag{4}$$

where $F(t)$ and μ are the cumulative distribution function and the mean of the intervals, respectively (Perkel et al., 1967). This is also the pdf for the duration from a random time to the previous firing.

Now, if there are 2 spike trains A and B which are independent, then the firings in train A are equivalent to random times with respect to train B, and vice versa. Thus, if A and B are independent, the pdf for the time from a spike in train A to the next spike in train B is

$$CI_{AB}(t) = \frac{1 - F_B(t)}{\mu_B}. \tag{5}$$

When independent, $CI_{AB}(t) = CI_{AB}(-t)$ (McFadden, 1962). Dependencies between A and B can then be detected by comparing the CI for positive and negative t, a significant difference indicating the presence of interaction. CI_{BA} is not, however, simply a scaled version of CI_{AB}. This can be seen when neuron A fires much more frequently than neuron B. Given a spike in A, the expected time until a firing in B will be relatively long; given a firing in B, the expected time until a firing A will be relatively short. Thus, for full cross-interval information both $CI_{AB}(t)$ and $CI_{BA}(t)$ must be calculated for positive t.

Cross-interval detection

Much of the problem with analyzing neural ensembles is simply the time and computational effort needed to analyze all possible interactions. Parametric techniques require complicated algorithms for estimating model parameters. Nonparametric techniques require pairwise cross-correlation analysis that increases in computational load proportionally to the square of the number of neurons in the ensemble. With cross-correlation analysis on an ensemble of n neurons using cross-correlation, predicted-correlation, residual-correlation and PST histograms, the number of histograms that are required for a complete analysis is $n(3n - 1)/2$. For example, an analysis of an ensemble of 8 neurons would require 92 histograms. This number could be significantly reduced, however, if there were a method which indicated, before the histograms were created, which pairs represent potential connectivity. Only those pairs would then need be analyzed. What is desirable is an algorithm which can be updated with each spike occurrence, using simple calculations, and thereby indicate at the end of the data

collection which neuron pairs are potentially connected.

Consider the cross-interval density function from neuron A to neuron B, CI_{AB}. With 2 neurons which are independent and firing at constant rate, CI_{AB} has the form $\lambda_B e^{-\lambda_B|t|}$ for both positive and negative t. This is also the form of the self-interval function, I_B. Johnson and Kiang (1976) used cross-intervals to estimate connectivity. After estimating λ from the number of spike occurrences, they subtracted the expected cross-interval function from the experimentally obtained function to create the cross-interval residual. If the residual function exceeded the confidence limits more than 5% of the time then a connection was detected. This method requires the computation of 2 cross-intervals at the end of data collection for every pair considered.

The mean duration for the cross-interval between A and B, as well as the self-interval for B, is $1/\lambda_B$. If neuron A excites neuron B such that the probability of B firing increases for a small duration after A fires, the mean of CI_{AB} will be shorter than the mean of I_B. It seems reasonable then to use the mean of the cross-intervals as a statistic for detection such that significant deviation from the self-interval mean would indicate a connection. This method is advantageous because it uses only the means of cross-intervals to make decisions and is thereby straightforward to implement.

A method which uses the mean of the cross-intervals to obtain an intensity measure will be derived in the next section. Three non-standard intensities will be used: the *excitation intensity*, λ_{ex}, which is the intensity of B during the excitation period; $\bar{\lambda}_{AB}(\Delta)$, the *mean intensity* of neuron B Δ seconds after A fires; λ_{AB}, the *relative intensity*, or intensity of B relative to A. These will be explained in detail later.

Analysis

For the sake of clarity, the only interaction examined in this section will be excitatory. Without loss of generality, the pre-synaptic neuron will always be neuron A and the post-synaptic neuron will always be neuron B. Also, the time t will always be with respect to the time after A has fired, i.e., the time axis in a cross-intensity function.

The model for interaction between neurons A and B will be as follows. The spontaneous firing rate of A and B will be defined as μ_A and μ_B. In the case of interaction from A to B, μ_B is the firing rate of B when A never fires. The excitatory connection will alter the firing rate of B such that

$$\mu_B(t) = \mu_B + \mu_{AB},\ t \leqslant \Delta$$
$$= \mu_B,\ t > \Delta. \tag{6}$$

The average intensities of the 2 neurons are then

$$\lambda_A = \mu_A$$
$$\lambda_B = \mu_B + \mu_{AB}\lambda_A\Delta \tag{7}$$

so that $\lambda_B = \mu_B$ if there is no connectivity or if the intensity of A is zero.

The excitatory effect is pulse-like since the intensity increases by μ_{AB} for Δ seconds after A fires. This is a simplification of effects seen in actual data, where the increase in intensity is more often exponentially increasing and decreasing. What is important in the modeling of excitatory effects by Eqn. 6 is that the area of the increased intensity in the model is the same area seen in actual data. This means that the increased probability of B firing after A fires is the same for both cases. The parameters that are adjusted to achieve this are μ_{AB} and Δ. It should also be noted that experimental results typically show a small latency period after the firing of the pre-synaptic neuron before the excitatory effect begins, while the model in Eqn. 6 has no such delay. This will be addressed in the Discussion section.

As previously stated, a method of detecting connectivity without calculating numerous histograms is desirable. With this in mind, the MLE of $\lambda_{ex} = \mu_B + \mu_{AB}$ will be derived from cross-interval data CI_{AB}, under the assumption that Δ is known. The excitation intensity, λ_{ex}, is the intensity of B during excitation from A and differs from λ_B only when an excitatory connection exists.

The cross-interval probability density function for a neuron described by Eqn. 6 is

$$CI_{AB}(t) = \lambda_{ex} e^{-\lambda_{ex} t},\ t \leqslant \Delta$$
$$= \lambda_B e^{-(\lambda_{ex}-\lambda_B)\Delta - \lambda_B t},\ t > \Delta. \quad (8)$$

This pdf corresponds to the intensity of neuron B changing from λ_{ex} to λ_B, Δ seconds after neuron A fires. The average intensity of neuron B does not change to its spontaneous rate μ_B after Δ seconds as in Eqn. 6 because there is a possibility that neuron A will fire again before B fires. Thus, there is a certain probability that the intensity for neuron B will be μ_B but there is also the probability that the intensity will be $\mu_B + \mu_{AB}$. The overall *average* intensity is then $\lambda_B = \mu_B + \lambda_A \mu_{AB} \Delta$.

The pdf for the N cross-intervals is

$$\left\{ \prod_{i=1}^{N_1} \lambda_{ex} e^{-\lambda_{ex} t_i} \right\} \left\{ \prod_{j=1}^{N_2} \lambda_B e^{-(\lambda_{ex}-\lambda_B)\Delta - \lambda_B t_j} \right\}, \quad (9)$$

where N_1 is the number of cross-intervals of duration less than or equal to Δ and N_2 is the number of cross-intervals of duration greater than Δ. This distribution will be maximized under λ_{ex} in order to obtain the MLE of λ_{ex}, assuming that Δ is known.

Near-perfect knowledge: excitation intensity

The only unknown assumed in Eqn. 9 is λ_{ex}. λ_B is considered known and is estimated by N_b/T, where T is the observation period and N_B is the number of spikes fired by B. When the derivative with respect to λ_{ex} of the *log* of Eqn. 9 is taken and then set to zero, the maximum likelihood estimator (MLE) is obtained:

$$\hat{\lambda}_{ex} = N_1 / \left(\sum_{i=1}^{N_1} t_i + N_2 \Delta \right). \quad (10)$$

$\hat{\lambda}_{ex}$ is the maximum likelihood estimate of the intensity of B during $t \in [0,\Delta)$ using cross-interval information. This estimate is then compared to λ_B and if it is significantly different then a connection is detected.

This estimator was implemented on a simulated 3-neuron ensemble with $\lambda_i = 100$ spikes/s, $\Delta = 1$ ms and $\mu_{ij} = 150$ spikes/s for the specified connections and 0 otherwise. Ten seconds of data were simulated using the point-process generator described by Snyder and Miller (1991).

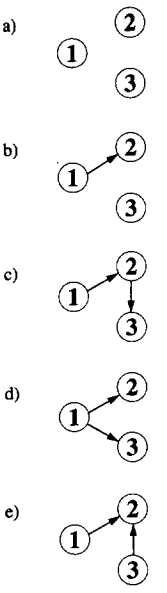

94.8	89.4	77.5
115.6	95.8	93.2
114.2	97.0	103.6

101.9	**257.7**	98.3
104.5	101.8	73.0
110.0	80.0	102.5

100.5	**237.6**	103.1
96.3	102.4	**261.8**
96.2	108.5	102.9

100.6	**311.2**	**281.1**
89.0	106.9	114.4
86.9	101.1	98.1

101.3	**239.2**	121.9
91.6	100.9	100.3
108.3	**246.3**	103.9

Fig. 2. Different ensemble configurations with excitatory connections. The component (i,j) in each matrix estimates the excitation intensity λ_{ex} using Eqn. 10, and the bold entries are the estimates for the actual connections. All connected pairs have $\lambda_{ex} = 250$ spikes/s; all unconnected have $\lambda_{ex} = 100$ spikes/s.

The results are given in the form of a matrix where entry (i,j) represents $\hat{\lambda}_{ex}^{ij}$, the excitation intensity estimated from CI_{ij}, and (i,i) represents the average firing rate λ_i estimated by N_i/T. Different interconnections were created to test the ability to distinguish between direct and indirect connectivity, and Fig. 2 shows both the neural ensembles and the matrices of estimators. The excitation intensity estimate $\hat{\lambda}_{ex}^{ij}$ in (i, j) must be compared with the average intensity λ_j in (j, j). The matrix values which represent connections are significantly different from the diagonal entries, signifying detection of a connection. Those values representing connectivity detection are bolded.

Fig. 2c shows that the estimator is not affected by indirect connectivity since (1,3) is not significantly different from (3,3). Common drivers do not significantly affect the estimator, as seen by

nearness of (2,3) to (3,3) and (3,2) to (2,2) in Fig. 2d.

The MLE of Eqn. 10 is a useful detector of connectivity since it requires minimum analysis of the data: the number of cross-intervals smaller and larger than Δ, and the sum (or average) of the cross-intervals small than Δ. The estimation matrix can be updated with every spike and requires the storage of 3 scalars for every potential connection analyzed.

One obvious fault with this method is that in reality the excitatory duration Δ is unknown and requires either a cross-interval or cross-correlation interval histogram to estimate. The maximum likelihood estimate of Δ is not simple and would require calculation from a histogram, a requirement which the technique presented here is intended to avoid. A priori information might be used, however, in the form of a maximum possible Δ. This will be explored in the next section.

Bounded Δ: mean intensity

The estimator of Eqn. 10 effectively compares B's intensity for $t \leq \Delta$ with the overall average intensity. In most experimental data, the excitatory effect occurs within a limited time after the pre-synaptic neuron fires. If one assumes that any excitatory effect will have died out 5 ms after the pre-synaptic neuron fires, Δ can be bounded by 5 ms. The estimator can then be effectively used to compare the average intensity for $t < 5$ ms with the overall average intensity. If there has been an increase in the intensity of any duration during the first 5 ms, then the estimated intensity during the first 5 ms will be greater than the average intensity.

When an estimate of Δ is used instead of the actual value, the MLE of Eqn. 10 estimates the average value of the intensity of $\hat{\Delta}$, even though there may be 2 distinct regions with different intensities during $(0,\hat{\Delta})$. This estimate of the average intensity, which is not only a function of the neural process but of the $\hat{\Delta}$ used, is termed the mean intensity and is denoted by $\bar{\lambda}(\hat{\Delta})$. Note that when $\hat{\Delta} = \Delta$, the actual duration of the excitation, $\bar{\lambda}(\Delta) = \lambda_{ex}$. The estimator of Eqn. 10 for $\bar{\lambda}_{ij}(\hat{\Delta})$ will be used with an estimate of $\hat{\Delta} = 5$ ms.

a)

95.8	94.8	93.1
102.7	95.8	96.2
102.9	101.9	103.6

94.8	94.0	93.9
97.2	95.8	100.8
107.0	108.6	103.6

b)

90.8	**123.0**	87.7
99.7	100.1	102.3
113.2	101.8	101.3

102.7	**119.0**	102.7
99.4	108.7	103.6
104.9	99.4	104.4

c)

105.6	**135.8**	110.5
91.7	97.9	**143.6**
96.4	90.7	101.3

106.2	**120.4**	106.5
101.4	100.4	**112.4**
100.5	98.0	104.9

d)

100.5	**133.2**	**140.2**
92.1	101.6	106.2
101.7	101.0	102.2

107.2	**117.0**	**116.9**
110.5	108.6	106.3
97.4	96.8	99.8

e)

100.1	**136.5**	87.0
105.2	103.0	97.7
103.8	**136.8**	101.1

101.6	**117.8**	102.9
90.7	98.6	92.7
102.7	**135.4**	104.6

I II

Fig. 3. I: the same as Fig. 2 only estimating the mean intensity using an estimate of $\Delta = 5$ ms instead of the true value of 1 ms in (10). II: the same as Fig. 2 only the component (i,j) in each matrix represents the estimate of the relative intensity λ_{ij}. The values bolded in both figures are the estimates of actual connections.

The left-hand column of Fig. 3 illustrates this technique using the same neural simulation as in Fig. 2, only $\hat{\Delta} = 5$ ms is used in Eqn. 10 instead of the true $\Delta = 1$ ms. The results are not as striking as in fig. 2 when Δ was known, but the connections can still be readily detected. The estimated values of $\bar{\lambda}_{ij}(\hat{\Delta})$ where there is a connection are not as large as before since the 1 ms excitatory intensity is spread out over 5 ms.

Unknown Δ: relative intensity

If one does not want to make any assumptions about Δ, the limit of Eqn. 10 can be taken as $\Delta \to \infty$ to obtain

$$\hat{\lambda}_{AB} = \lim_{\hat{\Delta} \to \infty} \bar{\lambda}(\hat{\Delta}) = N / \sum_{i=1}^{N} t_i. \tag{11}$$

This result is obtained because as $\Delta \to \infty$, $N_2 \to 0$ and $N_1 \to N$, the total number of cross-intervals. While letting $\Delta \to \infty$ is a naïve approach – it assumes the excitation after A fires lasts until B fires without decay — the estimator is in a very

simple form. It can be shown that the estimator is the MLE of λ assuming that the cross-interval has a pdf of the form $\lambda \exp(-\lambda t)$. The estimate $\hat{\lambda}_{AB}$ can be thought of as a *relative intensity*, i.e., the intensity of neuron B as seen by or relative to neuron A. If this relative intensity is different from the average intensity λ_B, then a connection can be said to exist.

An intuitive description and a justification of the estimator in Eqn. 11 can be seen when considering the Neyman–Pearson test (Van Trees, 1968). The Neyman–Pearson test is used when there are 2 hypotheses H_0 and H_1, where the statistics of H_0 are known but the statistics of H_1 are not. In this case, H_0 represents the hypothesis that there is no connection between neurons A and B, and H_1 represents the hypothesis that there is an excitatory connection from A to B. Under H_0, the distribution of the cross-intervals from A to B is

$$CI_{AB}(t) = \lambda_B e^{-\lambda_B t}. \tag{12}$$

Since λ_B is known, the statistics under H_0 are known.

Under H_1, the distribution of the cross-intervals is discontinuous and given in Eqn. 8. There are 2 independent unknowns in this equation: Δ and λ_{ex}. H_1 in this case is termed a composite hypothesis since the parameters Δ and λ_{ex} have a range of possible values. Even though Δ and λ_{ex} are unknown, however, they are not random variables but are constants. The subsequent lack of a probability density makes any Bayes test unreasonable, since a Bayes test would require probability distributions for Δ and λ_{ex} in order to create a likelihood ratio test. The Neyman–Pearson test, however, can be used as a method for detection since only the statistics under H_0 are used, i.e., the distributions of Δ and λ_{ex} are not needed since they exist only under H_1. Under the Neyman–Pearson test, the probability of a false alarm (PFA) is specified — which is definable since the statistics under H_0 are known — and bounds are derived. This creates decision regions for H_0 and H_1 based upon a significant statistic, which in this case is the relative intensity estimator of Eqn. 11. Under H_0, $\mu_{AB} = 0$ so the estimator $\hat{\lambda}_{AB}$ in Eqn. 11 estimates $\mu_B (= \lambda_B)$.

Appendix A derives an expression for γ, the upper bound for the decision region of H_0. It shows

$$\gamma = \frac{\lambda_B}{a''\sqrt{N}/(N-1) + N/(N-1)} \approx \frac{\lambda_B}{a''/\sqrt{N} + 1}, \tag{13}$$

where the approximation holds for large N. a'' is solely determined by the predetermined PFA and is defined by the value

$$\text{PFA} = 1 - erfc_*(a'') - \int_{a''}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right) du. \tag{14}$$

The lower limit a'' can be obtained from standard error function tables. Since a'' is negative for PFA < 0.5, γ is monotone decreased with respect to N and monotone increasing with respect to λ_B, which is not surprising since the expected value of $\hat{\lambda}_{AB}$ is λ_B under H_0.

Fig. 4 plots the bound γ as a function of N for PFAs of 0.05, 0.01 and 0.001. The bound is given as a multiple of λ_B since the form in Eqn. 13 can be easily rearranged into a ratio between γ and λ_B. For example, given a data set of 500 intervals with a firing rate of 100 spikes/s, the upper

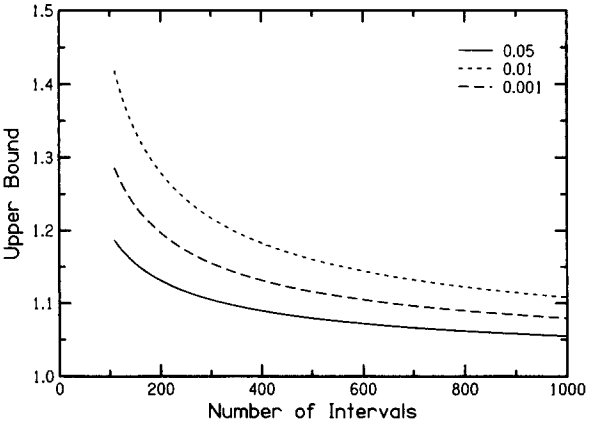

Fig. 4. The upper bound on the estimator Eqn. 11 for determining the hypothesis H_0 that no connection exists, given as a multiple of λ_B. The different lines represent different probabilities of false alarm.

bound for a PFA of 0.01 is approximately 111 spikes/s.

With γ set as determined by Eqn. 13, the detection characteristics are completely specified under H_0. To describe the detection characteristics under H_1, and thus completely characterize the detection task, the probability of detection (PD) must be derived. This is done in Appendix B, which results in

$$\begin{aligned} PD &= 1 - erfc_*\left[a'' + \sqrt{N}(\lambda_{ex} - \lambda_B)\Delta\right] \\ &= 1 - erfc_*\left[a'' + \sqrt{N}\,\mu_{AB}(1 - \lambda_A\Delta)\Delta\right]. \end{aligned} \tag{15}$$

PD is seen to be monotone increasing with Δ, N and μ_{AB}, as expected. PD is monotone decreasing with λ_A because an increase in λ_A decreases the difference between λ_{ex} and λ_B, which decreases PD as seen in the first line of Eqn. 15.

Fig. 5 plots the PD as a function of $(\lambda_{ex} - \lambda_B)\Delta$ and N for PFA = 0.05. $(\lambda_{ex} - \lambda_B)\Delta$ is termed the normalized efficacy, e'. The true efficacy value is $e = \mu_{AB}\Delta$ and, from Eqn. 15, $e' = (\lambda_{ex} - \lambda_B)\Delta = e(1 - \lambda_A\Delta)$. The difference value $(\lambda_{ex} - \lambda_B)\Delta$ is often mistakenly termed the efficacy because it is the area increase in the cross-correlation function due to connectivity. As explained in detail by Aertsen et al. (1989), this measure is not conditioned correctly. Fig. 5a plots PD versus e' with the number of intervals N as a parameter. Fig. 5b plots PD versus N with e' as a parameter.

In the simulation of networks shown in Fig. 2, the number of cross-intervals generated were approximately 1000 per pair and the normalized efficacy was 0.135. According to Fig. 5a the probability of excitation detection should be near 1. With the neurons firing at 100 spikes/s, the confidence limit for PFA = 0.05 is derived using Eqn. 13 as $\gamma = 109$ spikes/s. The results are shown in the right-hand column of Fig. 3 for the various ensemble connections, with the estimates that exceed the bound in bold. As can be seen, the detector was perfect, with no false alarms and no misses.

Discussion

The problem of detecting connectivity in an ensemble of neurons has been addressed with an emphasis on minimal computation time. The methods developed have used only cross-interval information and represent pairwise relationships with scalar values that can be updated with each spike occurrence. The simplicity of these methods can be see in Figs. 2 and 3, where a $n \times n$ matrix is used to represent possible pairwise connections.

All 3 methods involve an estimation of the excitation intensity. The first method, where the duration of the excitation Δ is known and λ_{ex} is estimated, represents near perfect knowledge of

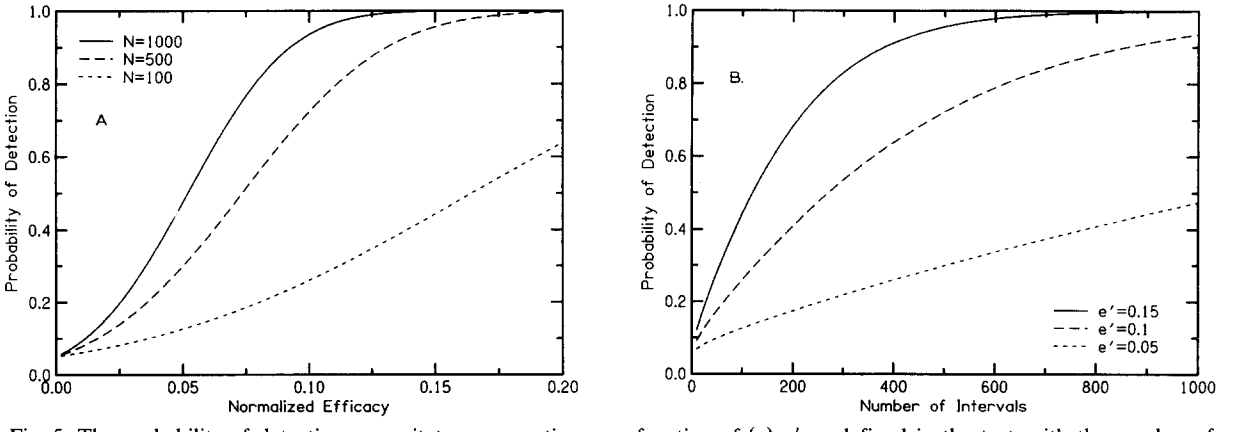

Fig. 5. The probability of detecting an excitatory connection as a function of (a) e' as defined in the text, with the number of cross-intervals N as the parameter; (b) the number of cross-intervals with e' as the parameter. Both plots are for PFA = 0.05.

the interaction. This is unrealistic since the timing characteristics of interactions vary and estimating the duration requires a complexity of analysis that these methods are attempting to avoid. The second method assumes that the region of excitation can be bounded, thereby upper bounding Δ. The mean intensity $\bar{\lambda}(\hat{\Delta})$ over the bounded region is estimated, effectively spreading the excitatory intensity over the bounded region such that $\bar{\lambda}(\hat{\Delta})$ is smaller than λ_{ex}. The third method makes no assumptions about the timing of the excitation and estimates the relative intensity λ_{AB}. The relative intensity is calculated by assuming that the cross-intervals are a result of a constant intensity and the intensity is estimated taking the inverse of the cross-interval mean. For excitatory interaction, the relative intensity λ_{AB} will be higher than the true intensity λ_B and the difference between the 2 will be greatest for high efficacy, or high μ_{AB} and Δ.

The differences between the 3 detectors are seen in Figs. 2 and 3. The perfect knowledge detector performs the best, as expected, with the intensities representing connections clearly distinguishable from the intensities with no connections. As stated previously, the relative intensity method is equivalent to the bounded method with the upper bound on Δ being ∞. The variance in the excitation intensity estimators decreases as $\hat{\Delta}$ increases because more cross-intervals N_1 are used to estimate the excitation intensity, as seen in Eqn. 10. The expected value of the estimate, however, decreases as $\hat{\Delta}$ increases because the excitation intensity is averaged across the bound. This is realized in Figs. 2 and 3 where the estimated intensities for connected pairs are smaller when $\hat{\Delta} = \infty$ than when $\hat{\Delta} = 5$ ms, and both are smaller than when $\hat{\Delta}$ is the true value of 1 ms.

One surprising feature of Figs. 2 and 3 is the absence of false alarms when there are indirect connections or common drivers. Fig. 2c shows a common situation where many techniques detect a connection between neurons 3 and 1 since neuron 3 is indirectly excited by neuron 1 through neuron 2. Entry (1,3) in all 3 detection matrices for this network show a negligible excitation intensity. Fig. 2d depicts a situation where neurons 2 and 3 could be falsely detected as connected since they are both timed to a common driver. Again, entries (2,3) and (3,2) show negligible excitation intensity.

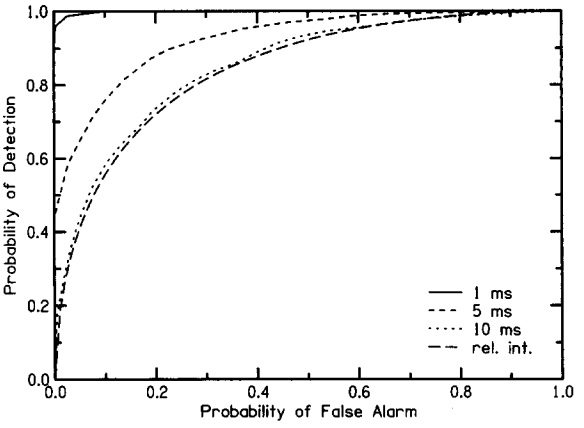

Fig. 6. Receiver operator characteristics for four different estimators, with $N = 1000$ cross-intervals, $\Delta = 1$ ms, $\lambda_i = 100$ spikes/s, and $\mu_{AB} = 50$ spikes/s. The solid line represents the excitation intensity estimator, the dashed line represents the mean intensity estimator with $\hat{\Delta} = 5$ ms, the dotted line represents the mean intensity estimator with $\hat{\Delta} = 10$ ms, the the long-dashed line represents the relative intensity estimator. The last one was derived using Eqn. 15, the rest were obtained using Monte Carlo simulations.

The intensities λ_{ex} and λ_{AB} are special cases of $\bar{\lambda}(\hat{\Delta})$, where $\lambda_{ex} = \bar{\lambda}(\Delta)$ and $\lambda_{AB} = \bar{\lambda}(\infty)$. The performance of $\bar{\lambda}(\hat{\Delta})$ for varying $\hat{\Delta}$ is shown in Fig. 6, where the receiver operator characteristics (ROC) are plotted for the different methods (Green and Swets, 1966). Both these values are a function of the detection threshold that is chosen, i.e., γ in Eqns. (A8) and (B5). The ROC curve for the relative intensity estimator was determined using Eqn. 15 while the rest were obtained using Monte Carlo simulations. The ROC plot was derived using the following characteristics: $N =$ 1000 cross-intervals, $\mu_{AB} = 50$ spikes/s, $\lambda_i = 100$ spikes/s and $\Delta = 1$ ms. These values represent approximately 10 s of data with a total of approximately 1000 spikes/neuron.

The lines represent $\hat{\Delta} = 1$ ms, 5 ms, 10 ms and ∞. The estimator for λ_{ex} is seen to perform nearly perfectly, with the sub-optimal region appearing in the top left corner. The performance of the mean intensity estimator with $\hat{\Delta} = 10$ ms is equivalent to the relative intensity estimator's

performance, indicating that λ_{AB} is a reasonable statistic to estimate when tight bounds on Δ cannot be used.

The performance of the relative intensity estimator is shown in Fig. 5 for PFA = 0.05. When making a decision about connectivity, correctly determining that connectivity exists is more important than correctly determining that it does not exist. Thus, increasing the probability of false alarm is a reasonable sacrifice in order to increase the probability of detection. Fig. 7 shows the same plots as in Fig. 5 for PFA = 0.25. The plots have shifted to the left, reducing the intensity difference and the number of intervals needed for an equivalent probability of detection.

Since the number of unconnected pairs is typically much greater than the number of connected pairs in an ensemble, the reduction in analysis afforded by using the relative intensity estimator is approximately to the value of the PFA. So, if a PFA of 0.25 is used, the analysis computations are reduced by 75%. In the original example of an 8 neuron ensemble, approximately 23 histograms would have to be analyzed instead of 92. Even with a PFA of 0.5, an unacceptable high value for most communication systems, the number of computations would be reduced by approximately one-half. In this case, the detection task reduces to choosing H_1 if $\lambda_{AB} > \lambda_B$ and H_0 if $\lambda_{Ab} \leq \lambda_B$.

The performance of the detector was seen in Eqn. 15 to become poorer with increasing firing rates. While many systems have neurons with rates on the order of 100 spikes/s (cochlear nucleus, olivary complex) others fire at a rate of just a few spikes/s (vestibular system, cortex). For low firing rates, the normalized efficacy discussed is approximately the true efficacy, i.e., $e' = \mu_{AB}\Delta(1 - \lambda_A\Delta) \approx \mu_{AB}\Delta = e$. This reduction in firing rate will reduce the number of cross-intervals N that are produced, however. Thus, there exists a trade-off.

One of the simplifying assumptions in the form of the connectivity, as seen in Eqn. 6, was that the excitation occurred instantly after the pre-synaptic neuron fired, i.e., there was no latency period before the post-synaptic neuron's intensity increased. This is not realistic since all experimental data has shown some form of latency period before the excitation occurred. For example, Michalski et al. (1983) have found that the average latency before excitation in the cat striate cortex was 0.62 ms, followed by an average excitation period of 2.22 ms. Part of the latency that is seen in cross-correlation functions can be attributed to separation effects due to estimating multiple-unit firings from a single electrode (Gochin et al., 1989), but a time delay between the pre-synaptic firing and the post-synaptic excitation does exist.

The effect of such a latency on these techniques would be to increase the cross-interval mean and decrease the estimated relative intensity. It can be shown that, while the expected

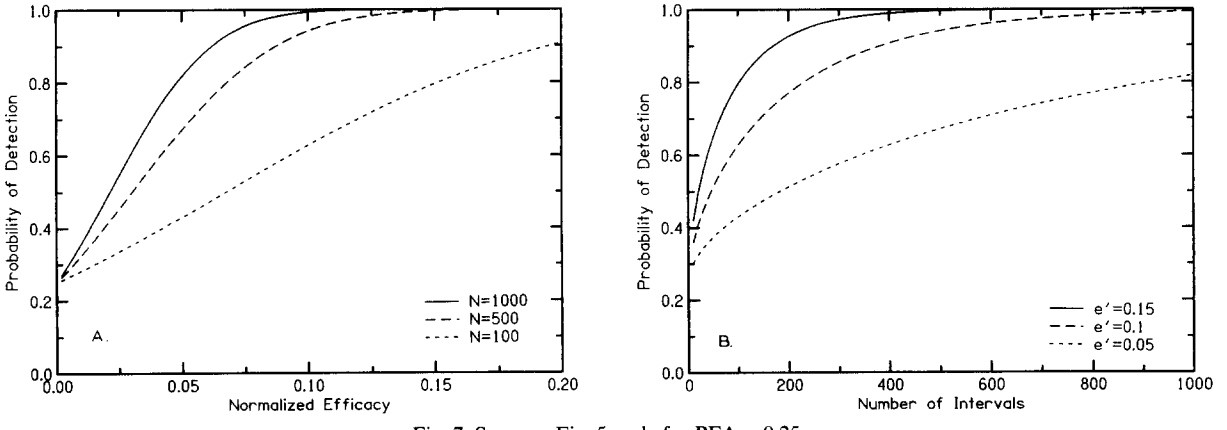

Fig. 7. Same as Fig. 5, only for PFA = 0.25.

relative intensity without latency is

$$\frac{\lambda_B}{1-(\lambda_{ex}-\lambda_B)\Delta} \tag{16}$$

the expected relative intensity with a latency of δ is

$$\frac{\lambda_B}{1-(1-\lambda_B\delta)(\lambda_{ex}-\lambda_B)\Delta}, \tag{17}$$

for small $\lambda_B\delta$ and $\lambda_{AB}\Delta$. The effect of δ on $\hat{\lambda}_{AB}$ is seen to be negligible for realistic parameter values.

Another assumption made in the analysis was that the neurons had independent firing rates of μ_i. This value can be either a spontaneous firing rate or a constant firing rate due to a constant driving function, such as a high frequency tone. While the performance characteristics derived in this text do not hold for time-varying driving functions, the techniques might still work due to the simple fact that connected units will tend to fire near each other and thus the cross-intervals will have a shorter mean than the self-intervals' mean. Several simulations using a sinusoidal intensity of varying frequency and modulation depth have resulted in similar cross-intensity performance as shown in Figs. 2 and 3. Also, units with common driving functions that are time delayed may be detected as interconnected since the stimulus will have the same effect on the cross-intervals that a connection would have. Shuffle, shift, or PST predictors will have to be used after connections have been detected to eliminate the stimulus-induced connections.

The excitation and mean intensity estimators are similar to one created from a cross-correlogram approach. Coincidence counts could be created using a window the size of the bounded excitation $\hat{\Delta}$ and the cross-correlation intensity compared to the firing rate. The cross-correlogram does not provide for a computationally simple technique where the excitation duration is not bounded, however, and thus a computationally efficient detection technique without a bound on Δ is unique to the cross-interval method.

Finally, the analysis presented has been with respect to excitatory connections. Inhibitory connections can easily be addressed by making μ_{AB} in Eqn. 6 negative and the rest of the analysis will remain the same except that γ will be a lower bound. If one wants to apply this method to detect both excitatory and inhibitory connections, Eqns. (A4), (A8) and (B5) will have a lower bound in the integration along with the upper bound γ. A corresponding upper bound for the normalized gaussian must be calculated along with a'' in (A9) to determine the PFA. An excitatory connection will then be detected if the relative intensity λ_{AB} (or mean intensity $\bar{\lambda}(\hat{\Delta})$ if the time of the effect is bounded by $\hat{\Delta}$) is above the upper bound and an inhibitory effect will be detected if λ_{AB} is below the lower bound.

Acknowledgements

This research was supported by a grant from the NIH (NIH NS21440). The authors would like to thank Dr. David J. Anderson and Dr. Ben M. Clopton for their helpful comments on this research. The authors would also like to thank the anonymous reviewer for their helpful comments on the original manuscript.

Appendix A

We will derive here the upper bound for the H_0 decision region. The pdf of the sum of N Poisson distributed intervals, each with pdf$\lambda e^{-\lambda t}$, is

$$p_{\Sigma t}(z) = \frac{(\lambda z)^{N-1}}{(N-1)!}\lambda e^{-\lambda z}. \tag{A1}$$

The MLE for λ given this distribution is $N/\Sigma_1^N t_i$. Letting $z = \Sigma_1^N t_i$, the expected value of N/z is

$$\begin{aligned} E[N/z] &= \int_0^\infty \frac{N(\lambda z)^{N-1}}{z(N-1)!}\lambda e^{-\lambda z}dz \\ &= \frac{N\lambda}{N-1}\int_0^\infty \frac{(\lambda z)^{N-2}}{(N-2)!}\lambda e^{-\lambda z}dz \\ &= \frac{N\lambda}{N-1}. \end{aligned} \tag{A2}$$

Thus, the MLE of λ_{AB} from the cross-interval data is biased by a factor of N/(N − 1), and the unbiased estimator of λ_{AB} is

$$\hat{\lambda}_{AB} = (N-1)/\sum_{i=1}^{N} t_i. \qquad (A3)$$

The pdf for $\Sigma_1^N t_i$ under H_0 is given by (A1). Defining the variable $y = (N-1)/\Sigma_1^N t_i$, the upper bound γ must be determined such that

$$1 - \int_0^{\gamma} p_y(Y)\,dy = PFA, \qquad (A4)$$

where again PFA is the probability of a false alarm, that is, the probability of deciding H_1 when the signal H_0.

Unfortunately, $p_y(Y)$ is difficult to integrate to a closed form solution for arbitrary bounds. Since $z = \Sigma_1^N t_i/N$ is the normalized sum of identically distributed variables, however, the pdf for z is in the shape of a gaussian. Using the Central Limit Theorem (Feller, 1968),

$$p_z^{H0}(Z) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left[-\frac{1}{2}\left[\frac{(Z-m_0)}{\sigma}\right]^2\right], \qquad (A5)$$

where $m_0 = \mathrm{E}[z] = 1/\lambda_B$ and $\sigma^2 = var[z] = 1/(N\lambda_B^2)$. The integral of (A5) over [0,∞) is assumed to be 1, i.e., the area of the pdf over negative values of z is negligible. This is justifiable because the actual area over negative values is $0.5 - \mathrm{erf}(\mu/\sigma) = 0.5 - \mathrm{erf}(\sqrt{N})$ which is negligible for large N.

Recalling that the estimator of interest is $(N-1)/\Sigma_1^N t_i$, define $y = (N-1)/\Sigma_1^N t_i = (N-1)/(Nz)$. When a variable y is a function of z, the pdf for y, $p_y(Y)$ can be derived from $p_z(Z)$ using

$$p_y(Y) = p_z(Z)\left|\frac{dz}{dy}\right|. \qquad (A6)$$

From (A5) and (A6),

$$p_y(Y) = \frac{N-1}{Ny^2}\frac{1}{\sqrt{2\pi}\sigma} \times \exp\left[-\frac{1}{2}\left(\frac{(N-1)/(Ny) - \mu}{\sigma}\right)^2\right]. \qquad (A7)$$

To derive the Neyman–Pearson test, the bound γ must be determined for which the integration of $p_y(Y)$ over $[\gamma,\infty)$ equals the predetermined PFA. Thus, γ must be determined such that

$$\begin{aligned} 1 - PFA &= \frac{N-1}{n}\int_0^{\gamma}\frac{1}{y^2}\frac{1}{\sqrt{2\pi}\sigma} \\ &\quad \times \exp\left[-\frac{1}{2}\left(\frac{(N-1)/(Ny) - m_0}{\sigma}\right)^2\right]dy, \\ &\quad 0 < y < \infty \\ &= \int_{a'}^{\infty}\frac{1}{\sqrt{2\pi}\sigma}\exp\left[-\frac{1}{2}\left(\frac{x-m_0}{\sigma}\right)^2\right]dx, \\ &\quad 0 < x < \infty \\ &= \int_{a''}^{\infty}\frac{1}{\sqrt{2\pi}}\exp\left[-\frac{u^2}{2}\right]du, \\ &\quad -\frac{m_0}{\sigma} < u < \infty, \end{aligned} \qquad (A8)$$

where the bounds are

$$a' = \frac{N-1}{N\gamma}$$

$$a'' = \frac{(N-1)/(N\gamma) - m_0}{\sigma}. \qquad (A9)$$

In (A8), the second line was obtained by substituting $x = (N-1)/(Ny)$ and the third line was obtained by substituting $u = (x - m_0)/\sigma$. The function in the third line is the normalized gaussian and the bound a'' is solely determined by PFA and can be calculated using standard error function tables. The original bound γ can be derived from a″ by rearranging Eqn. (A9):

$$\begin{aligned} \gamma &= \frac{N-1}{N(a''\sigma + m_0)} \\ &= \frac{\lambda_B}{a''\sqrt{N}/(N-1) + N/(N-1)} \\ &\approx \frac{\lambda_B}{a''/\sqrt{N} + 1}. \end{aligned} \qquad (A10)$$

Appendix B

Here will we derive an expression for PD given a specified PFA. The expected value of an interval with a pdf described by Eqn. 8 is

$$\begin{aligned} \mathrm{E}[t_i] &= 1/\lambda_{ex} + \exp(-\lambda_{ex}\Delta)(1/\lambda_B - 1/\lambda_{ex}) \\ &\approx (1-\lambda_{ex}\Delta)/\lambda_B + \Delta. \end{aligned} \quad (B1)$$

The variance can be shown to be approximately

$$\begin{aligned} \mathrm{var}[t_i] &\approx (1-\lambda_{ex}\Delta)(1+\lambda_{ex}\Delta)/\lambda_B + \Delta^2, \\ &\quad \mu_A\Delta \ll 1 \\ &\approx 1/\lambda_B^2,\ \lambda_{ex}\Delta \ll 1. \end{aligned} \quad (B2)$$

Using the same method applied in Appendix A, the pdf for $z = \sum_1^N t_i/N$ is

$$p_z^{H1}(Z) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left[-\frac{1}{2}\left(\frac{Z-m_1}{\sigma}\right)^2\right], \quad (B3)$$

with

$$\begin{aligned} m_1 &= (1-\lambda_{ex}\Delta)/(N\lambda_B) + \Delta/N \\ \sigma^2 &= 1/(N\lambda_B^2). \end{aligned} \quad (B4)$$

The probability of detection PD is defined as

$$\begin{aligned} 1-PD &= \frac{N-1}{N}\int_0^{\gamma}\frac{N-1}{y^2}\frac{1}{\sqrt{2\pi}\sigma} \\ &\quad \times\exp\left[-\frac{1}{2}\left(\frac{(N-1)/(Ny)-m_1}{\sigma}\right)^2\right]dy \\ &= \int_{a'}^{\infty}\frac{1}{\sqrt{2\pi}\sigma}\exp\left[-\frac{1}{2}\left(\frac{x-m_1}{\sigma}\right)^2\right]du \\ &= \int_{a''}^{\infty}\frac{1}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\left(u-\frac{m_1-m_0}{\sigma}\right)^2\right]du, \end{aligned} \quad (B5)$$

where the same variable substitutions were made as in (A8) and thus the bounds a′ and a″ are also the same as in (A9), with m_1 instead of m_0. Comparing (A8) with (B5), the 2 hypotheses H_0 and H_1 can be modeled as 2 gaussian distributions with unity variance and a mean separation d of $(m_1 - m_0)/\sigma$. The separation can be reduced to

$$\begin{aligned} d &= \frac{m_1-m_0}{\sigma} \\ &= \frac{(1-\lambda_{ex}\Delta)/\lambda_B + \Delta - 1/\lambda_B}{1/(\sqrt{N}\lambda_B)} \\ &= \sqrt{N}(\lambda_B - \lambda_{ex})\Delta. \end{aligned} \quad (B6)$$

The ideal threshold for such a case is $a'' = d/2$, but the threshold was already predefined for a specific PFA in Eqn. 14.

With a'' defined for a specific PFA, the PD can be expressed from (B5) as

$$\begin{aligned} \mathrm{PD} &= 1 - erfc_*[a'' - (m_1-m_0)/\sigma] \\ &= 1 - erfc_*[a'' + \sqrt{N}(\lambda_{ex}-\lambda_B)\Delta] \\ &= 1 - erfc_*[a'' + \sqrt{N}(\mu_B + \mu_{AB} - \mu_B \\ &\quad -\lambda_A\mu_{AB}\Delta)\Delta] \\ &= 1 - erfc_*[a'' + \sqrt{N}\,\mu_{AB}(1-\lambda_A\Delta)\Delta], \end{aligned} \quad (B7)$$

where

$$erfc_*(x) = \int_x^{\infty}\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{u^2}{2}\right)du. \quad (B8)$$

References

Aertsen, A.M.H.J. and Gerstein, G.L. (1985) Evaluation of neuronal connectivity: sensitivity of cross-correlation. Brain Res., 340: 341–354.

Aersten, A.M.H.J., Gerstein, G.L., Habib, M.K. and Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of "effective connectivity", J. Neurophys., 61: 900–918.

BeMent, S.L., Wise, K.D., Anderson, D.J., Najafi, K. and Drake, K.L. (1986) Solid-state electrodes for multichannel multiplexed intracortical neuronal recording. IEEE Trans. Biomed. Eng., BME-33: 230–241.

Blum, N.A., Carkhuff, B.G., Charles, Jr., H.K., Edwards, R.L. and Meyer, R.A. (1991) Multisite microprobes for neural recordings. IEEE Trans. Biomed. Eng., BME-38: 68–74.

Brillinger, D.R. (1988) Maximum likelihood analysis of spike trains of interacting nerve cells. Biol. Cybern., 59: 189–200.

Chornoboy, E.S., Schramm, L.P. and Karr, A.F. (1988) Maximum likelihood identification of neural point process systems. Biol. Cybern., 59: 265–275.

Edwards, B.W. and Wakefield, G.H. (1990) On the statistics of binned neural point processes: the Bernoulli approximation and AR representation of the PST histogram. Biol. Cybern., 64: 145–153.

Feller, W. (1968) An Introduction to Probability Theory and its Applications. John Wiley, New York.

Gerstein, G.L. and Perkel, D.H. (1972) Mutual temporal relationships among neuronal spike trains. Biophys. J., 12: 453–473.

Gerstein, G.L., Perkel, D.H. and Dayhoff, J.E. (1985) Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. J. Neurosci., 5: 881–889.

Gochin, P.M., Kaltenbach, J.A. and Gerstein G.L. (1989) Coordinated activity of neuron pairs in anesthetized rat dorsal cochlear nucleus. Brain Res., 497: 1–11.

Green, D.M. and Swets, J.A. (1974) Signal Detection Theory and Psychophysics. Krieger, New York.

Johnson, D.H. and Kiang, Y.S.K. (1976) Analysis of discharges recorded simultaneously from pairs of auditory nerve fibers. Biophys. J., 16: 719–734.

Johnson, D.H. and Swami, A. (1983) The transmission of signals by auditory-nerve fiber discharge patterns. J. Acoust. Soc. Am., 74(2): 493–501.

McFadden, J.A. (1962). On the lengths of intervals in a stationary point process. J. Roy. Statist. Soc., Ser. B, 24: 364–382.

Melssen, W.J. and Epping, W.J.M. (1987) Detection and estimation of neuronal connectivity based on crosscorrelation analysis. Biol. Cybern., 57: 403–414.

Michalski, A., Gerstein, G.L., Czarkowska, J. and Tarnecki, R. (1983) Interactions between cat striate cortex neurons. Exp. Brain Res., 51: 97–107.

Palm, G., Aersten, A.M.H.G. and Gerstein, G.L. (1988) On the significance of correlations among neuronal spike trains. Biol. Cybern., 59: 1–11.

Perkel, D.H., Gerstein, G.L. and Moore, G.P. (1967) Neuronal spike trains and stochastic point processes. II. simultaneous spike trains. Biophys. J., 7: 419–440.

Perkel, D.H., Gerstein, G.L., Smith, M.S. and Tatton, W.G. (1975) Nerve-impulse patterns: a quantitative display technique for three neurons. Brain Res., 100: 271–296.

Snyder, D.L. and Miller, M.I. (1991) Random Point Processes in Time and Space. Springer, New York.

Van Trees H.L. (1968) Detection, Estimation, and Modulation Theory. John Wiley, New York.

Voigt, H.F. and Young, E.D. (1990) Cross-correlation analysis of inhibitory interactions in dorsal cochlear nucleus. J. Neurophys., 64: 1590–1610.