Detecting unitary events without discretization of time

Sonja Grün [a,*], Markus Diesmann [b], Franck Grammont [c], Alexa Riehle [c], Ad Aertsen [b]

[a] *Department of Neurophysiology, Max-Planck-Institute for Brain Research, 60528 Frankfurt/Main, Germany*
[b] *Department of Neurobiology and Biophysics, Albert-Ludwigs University, Freiburg, Germany*
[c] *Center for Research in Cognitive Neuroscience, CRNC-CNRS, Marseille, France*

Abstract

In earlier studies we developed the 'Unitary Events' analysis (Grün S. Unitary Joint-Events in Multiple-Neuron Spiking Activity: Detection, Significance and Interpretation. Reihe Physik, Band 60. Thun, Frankfurt/Main: Verlag Harri Deutsch, 1996.) to detect the presence of conspicuous spike coincidences in multiple single unit recordings and to evaluate their statistical significance. The method enabled us to study the relation between spike synchronization and behavioral events (Riehle A, Grün S, Diesmann M, Aertsen A. Spike synchronization and rate modulation differentially involved in motor cortical function. Science 1997;278:1950–1953.). There is recent experimental evidence that the timing accuracy of coincident spiking events, which might be relevant for higher brain function, may be in the range of 1–5 ms. To detect coincidences on that time scale, we sectioned the observation interval into short disjunct time slices ('bins'). Unitary Events analysis of this discretized process demonstrated that coincident events can indeed be reliably detected. However, the method looses sensitivity for higher temporal jitter of the events constituting the coincidences (Grün S. Unitary Joint-Events in Multiple-Neuron Spiking Activity: Detection, Significance and Interpretation. Reihe Physik, Band 60. Thun, Frankfurt/Main: Verlag Harri Deutsch, 1996.). Here we present a new approach, the 'multiple shift' method (MS), which overcomes the need for binning and treats the data in their (original) high time resolution (typically 1 ms, or better). Technically, coincidences are detected by shifting the spike trains against each other over the range of allowed coincidence width and integrating the number of exact coincidences (on the time resolution of the data) over all shifts. We found that the new method enhances the sensitivity for coincidences with temporal jitter. Both methods are outlined and compared on the basis of their analytical description and their application on simulated data. The performance on experimental data is illustrated. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Neuronal interaction; Cell assembly; Multiple-neuron activity; Trial-by-trial detection; Synchronicity; Near-coincidences; Statistical analysis

1. Introduction

1.1. Detecting excess coincidences as Unitary Events

It is now generally accepted that both perceptual and motor functions are based on joint processing in neuronal networks which are widely distributed over various brain structures. However, it is much less clear, how these networks organize dynamically in space and time to cope with momentary computational demands. The concept emerged that computational processes in the brain could rely on the relative timing of spike discharges among neurons within such functional groups (von der Malsburg, 1981; Abeles, 1982, 1991; Gerstein et al., 1989; Palm, 1990; Singer, 1993), commonly called cell assemblies (Hebb, 1949). In this view, changes of the cooperative interplay among neurons within an assembly, induced by sensory and behavioral events, should be reflected in systematic and rapid modulations of precise timing of spike occurrences in the participating neurons. An essential ingredient of the notion of coordinated ensemble activity is its flexibility and dynamic nature. To critically test if such a temporal scheme is actually implemented in the central nervous system, it is necessary to simultaneously observe the activities of many neurons, and to analyze these activities for signs of temporal coordination. The op-

* Corresponding author. Tel.: +49-69-96769-336; fax: +49-69-96769-327.

E-mail address: gruen@mpih-frankfurt.mpg.de (S. Grün)

portunity to decipher the functional cooperativity among neurons was entranced by the recent development of new technologies for recording multiple single-neuron activities in brain structures of behaving animals. Associated with this development, new computational tools were designed to analyze and interpret the large amount of information in such multichannel recordings.

In the conceptual framework of distributed networks, it is particularly intriguing to trace the temporal evolution of cooperative neuronal activity within such networks. For that purpose, the Joint-peri-stimulus-time histogram (JPSTH, Aertsen et al., 1989), Gravitational Clustering (Gerstein and Aertsen, 1985), and various flavours of Hidden Markov Models (HMM, Abeles et al., 1995; Gat et al., 1997) have been developed. However, although the dynamics of synchronicity can be observed as a function of time by averaging over trials, it has so far not been possible to analyze individual spike coincidences on a trial by trial basis. With this goal in mind, we recently developed the 'Unitary Events' analysis (Grün, 1996; Grün and Aertsen, 1999a,b) for detecting the presence of conspicuous spike coincidences in multiple single neuron recordings and evaluating their statistical significance. Basically, this technique allows one to determine those spike coincidences which violate the assumption of independence of the participating neurons and insofar are an expression of the activation of a functional cell assembly (Aertsen et al., 1991). The statistical null-hypothesis is formulated on the basis of the individual firing probabilities of the participating neurons. By means of this null-hypothesis, it is possible to calculate the number of expected coincidences. As a result of calculating the statistical significance of the difference between expected and measured coincidences, one obtains both the amount and the moment in time of the significant excess coincident spiking activities ('Unitary Events'; for technical details, see Appendix B). To account for the dynamics of synchronized activity as well as to deal with non-stationarities in the firing rate of the neurons, synchronicity is estimated on the basis of small time segments, by sliding a boxcar window in steps along the data. This technique allows one to describe a detailed relationship between spike synchronization, rate variations and behaviorally relevant events (Riehle et al., 1997). Effectively, UE-analysis is strongly related to evaluating the dynamics and significance of the diagonal trace of the JPSTH-matrix (Aertsen et al., 1989). Also the significance measure used (the modified surprise function) is very similar. UE-analysis deviates from JPSTH-analysis, however, in that it is not satisfied with detecting significant dynamic correlation per se, but makes a first step towards recovering the actual events that constitute this dynamic correlation.

The usual time resolution of the data aquisition in electrophysiological recordings is less then or equal to 1 ms. There is recent experimental evidence that the timing accuracy of spikes, which might be relevant for higher brain functions, can be as precise as 1–5 ms (Abeles et al., 1993; Riehle et al., 1997). To detect synchronous spikes on a particular time scale, we sectioned the observation interval into short disjunct time slices ('bins') (disjunct binning, DB). After such binning, binary processes were constructed from each spike train by assigning a '1' to time slices in which one or more spikes occurred ('clipping') and '0' to time slices in which no spike occurred. Although coincident spiking events can reliably be detected by using such discretized process, the method looses sensitivity for higher temporal jitter of the coincident events (Grün, 1996). This is mainly due to the non-linear effect of binning and clipping of the single spike trains, on the one hand, and the application of the same binning grid over multiple spike trains, on the other.

Here we present an alternative approach, the 'multiple shift' method (MS). This method overcomes the need for binning, and thereby treats the data in their (original) high time resolution. Technically, coincidences are detected by shifting the spike trains against each other over the range of allowed coincidence width and integrating the number of exact coincidences (on the time resolution of the data) over all shifts.

We first present the analytical descriptions for both methods (a list of symbols used is given in Appendix A). We then compare the two methods using surrogate data sets. Conceptually, we separate spike trains in 'background' spikes, i.e. uncorrelated spikes, and spikes being involved in coincidences. Thus, in our simulations we first generate independent spike trains with a given background rate, and then 'inject' coincident spikes of a given coincidence width (tolerance) into both trains. The firing rate levels are chosen in physiologically plausible ranges. Our analytical descriptions are constrained to low coincidence rate levels, such that interactions of injected coincidences are neglectable.

In a second step, we compare the two methods for their reliability to detect near-coincidences. It turns out, that MS is more sensitive to detect near-coincidences than DB. To illustrate the performance of our method, we apply MS to a particular experimental data set. Based on the result of this analysis, we set up a simple model for the composition of the coincident spiking activity. Using our analytical description, we estimated the parameters of our model from the experimental result and verified our assumptions by simulations.

For simplicity, we only discuss two parallel processes, but this work serves as a basis for an expansion to M parallel processes.

2. Simulation experiment

In order to calibrate and test the analysis methods, we used simulated spike trains in which we could control the firing rates of the 'neurons' and the temporal precision of the spike coincidences (the coincidence width). Here, firing rates are composed of both background activity and coincident activity (see Fig. 1).

In a first step, spike trains of time duration T were generated independently as Poisson spike trains, simulating uncorrelated background activity. At each instant of time a random number x from the interval [0, 1] (equally distributed) was drawn and compared to the instantaneous firing probability p_r given by the background rate λ_r ($p_r = \lambda_r \cdot h$, see e.g. Papoulis (1991)). If $x \leq p_r$, a spike ('1') was set, otherwise not ('0'). The resulting spike trains were then put together in parallel, as if they were recorded simultaneously.

Synchronized activity was generated as a single spike train, together with a 'noisy' replica of it (same duration T and time resolution h as for the background spike trains). The master train was generated again as a Poisson process, with coincidence rate $\lambda_c (\ll \lambda_r)$. The noisy replica was generated by jittering each spike independently and uniformly over a time window of $\pm s$ time steps around its original position, thus simulating a coincidence width of s. The probability to be placed at any one of the $2s + 1$ possible positions is $p_s = 1/(2s + 1)$. Thus, the coincidence width could be controlled by the variable s.

Finally, synchronized activity was injected into the spike trains representing the background activity by merging the two sets of spike trains. If by chance, a 'coincidence' spike coincided with background spike, it was discarded, thereby effectively clipping that particular bin content to 1. The probability per time step for such clipping is given by the product: $p_r \cdot p_c$.

The resulting total firing probability per spike train is given by the sum of the probabilities for background spikes and the coincident spikes, reduced by the probability for clipped spikes:

$$p_{rc} = p_r + p_c - p_r \cdot p_c \tag{1}$$

3. Disjunct binning (DB)

3.1. Detection of near-coincidences

To enable our analysis to detect coincident spike events in simultaneously recorded spike trains with a certain tolerance regarding coincidence precision, we generate a new process on a less restrictive time scale by sectioning the observation interval T into

$$N(b) = \frac{T}{b \cdot h} \tag{2}$$

disjunct time segments (bins) of width b (in units of the original time resolution h). In order to treat this process as a binary process (possible outcomes within a bin $\epsilon \{0, 1\}$) data are clipped to 1 in the case of more than one spike within a bin.

Coincidences of these newly generated processes are detected on the new time scale $b \cdot h$. Simultaneous 1's in both spike trains are counted as a coincidence. The total count of all coincidences in all bins $N(b)$ provides the number of measured coincidences n.

The expected number of coincidences ($\hat{n}$) is calculated on the basis of the firing probabilities (p_1, p_2) of the individual (new) processes involved:

$$\hat{n} = p_1 \cdot p_2 \cdot N(b) \tag{3}$$

For evaluating the significance of the difference between the numbers of observed and expected coinci-
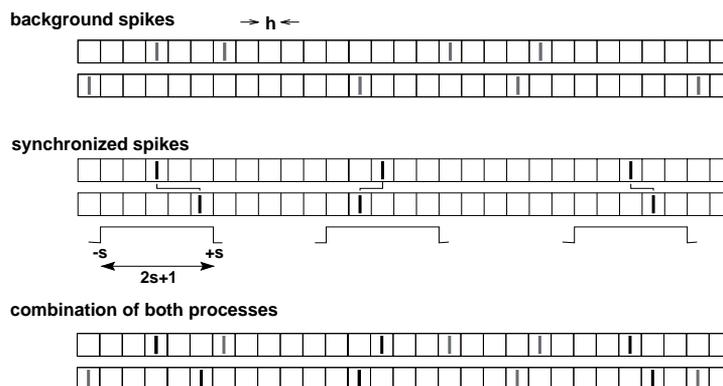
Fig. 1. Simulation experiment. Spike trains are composed of two independent contributions: background and coincident activity. Independently generated spike trains, generated as Poisson spike trains with given stationary rate λ_r, constitute the uncorrelated background activity. Synchronized activity is generated separately as a Poisson spike train with coincidence rate λ_c, and a jittered replica of it, by 'jittering' each spike relative to the original position by $\pm s$ time steps. The synchronized activity is then 'injected' into the background activity by adding the two activities time step by time step. If two spikes or more fall within a single time step, the contents is clipped to 1.
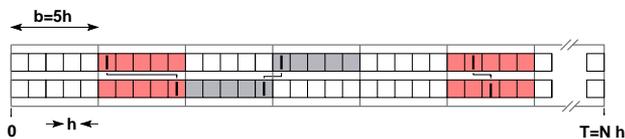
Fig. 2. Disjunct Binning. In order to detect near-coincidences of spike distances up to $s = b - 1$, the spike trains are sectioned in parallel into disjunct bins of width $b = s + 1$ (in units of h). Bins that contain a spike are marked in grey. Note, that two spikes in the middle, although they are less then $5h$ apart, are split by a bin border and thus cannot be detected as a coincident event.
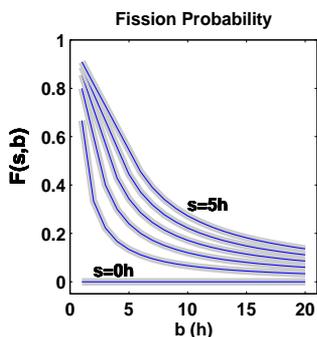
Fig. 3. Probability for splitting coincidences due to the binning procedure. The curves show the fission probability as a function of the binsize b (in units of the time resolution h) for a range of coincidence widths s of the injected coincidences. In black the fission probability calculated using Eq. (5), in grey using Eq. (4) for the series representation.

dences the joint-surprise may be used (Eq. (27) in Appendix B). The task now is to find analytical descriptions for n and $\hat{n}$ in our model situation (Section 2).

3.2. Analytical description: observed coincidences

The number of observed coincidences has two contributions: injected coincidences and coincidences obtained by chance due to the uncorrelated background activity. In a first step, the influence of the binning procedure on the injected coincidences will be examined.

3.2.1. Injected coincidences: splitting by binning

The application of the fixed binning grid to near-coincident events with a coincidence width s, that were injected at random instances in time, leads to situations in which two near-coincident events may fall into neighboring bins. Such near-coincidences are split and cannot be detected anymore (see Fig. 2). The fraction of split coincidences can be determined as follows.

Consider a spike in the first train at an arbitrary position. The probability for that spike to be at distance i from the nearest left bin border is: $P_i = N(b)/N = 1/b$. The probability for the second (near-coincident) spike to be at the left of that bin

border is: $P_l^i = (s - i) \cdot 1/(2s + 1)$. Thus, the total probability for near-coincident spikes to be separated by the left bin border is:

$$P_l = \sum_{i=0}^{r} P_i \cdot P_l^i, \quad \text{with } r = \begin{cases} s - 1 & \text{for } b \geq s \\ b - 1 & \text{for } b < s \end{cases} \quad (4)$$

By symmetry it follows that the total probability for near-coincident spikes to be splitted by a bin border equals $F(s, b) = 2 \cdot P_l$. The latter can be expressed in an arithmetic series representation, yielding an expression for the 'fission' probability as the result of binning:

$$F(s, b) = \begin{cases} \frac{1}{2s+1} \cdot \frac{s}{b} \cdot (s + 1) & \text{for } b \geq s \\ \frac{1}{2s+1} \cdot (2s - b + 1) & \text{for } b < s \end{cases} \quad (5)$$

The shape of $F(s, b)$ as a function of b is shown in Fig. 3 for different choices of s. As expected, exact coincidences ($s = 0$) are not affected since they cannot be splitted. For increasing coincidence widths s, the probability for splitting coincidences increases. It is high for small bin sizes, and decreases for larger bin sizes, finally converging to 0 for the case of one bin covering the whole data piece (Fig. 3). As a result, the total number of coincidences after binning is reduced by a factor $1 - F(s, b)$:

$$n_c(s, b) = p_c \cdot (1 - F(s, b)) \cdot N(b) \quad (6)$$

Note, that for a coincidence width of $s = 2$ and a binning grid of $b = 2$, for instance, the original coincidence probability is reduced by a factor 0.5. Thus, one fails to detect a considerable fraction of coincidences. For bin sizes smaller than the coincidence width s all coincidences with inter-spike distances larger or equal than the bin width are split. Only coincidences with a spike distance smaller than the bin width can be detected at all. For bin sizes larger than s, the near-coincidences are in principle detectable since they would fit into a single bin. However, for both cases it holds that some coincidences will be split due to their positioning relative to the binning grid.

3.2.2. 'Background' coincidences

The second contribution to the observed number of coincidences are the coincidences occurring by chance due to uncorrelated background activity. In addition to the regular background activity (p_r), those spikes have to be considered as background, that were originally part of a near-coincidence, but were split by the binning grid and now contribute to the background. This gives an additional contribution $p_c \cdot F(s, b)$ to the background rate. However, some of these debris may have collided with regular background spikes and thus have been clipped by the injection procedure. This results in a reduction of the background probability by the amount $p_r \cdot p_c \cdot F(s, b)$. Thus, for the total background spike probability per time step we obtain:

$$p'_r(s, b) = \underbrace{p_r}_{\text{backgr.}} + \underbrace{p_c \cdot F(s, b)}_{\text{debris}} - \underbrace{p_r \cdot p_c \cdot F(s, b)}_{\text{collisions}} \qquad (7)$$

This yields (see Appendix C) for the spike probability after binning and clipping:

$$p''_r(s, b) = 1 - (1 - p'_r(s, b))^b \qquad (8)$$

Assuming both neurons to have the same background rates, then the expected joint probability is obtained according to Eq. (3) by squaring:

$$\hat{P}_r = (p''_r(s, b))^2 \qquad (9)$$

The number of coincidences due to the background activity equals $\hat{P}_r$ multiplied with the number of available bins. The latter are the number of bins $N(b)$ reduced by the number of bins which are already 'occupied' by injected coincidences (i.e. $n_c(s, b)$):

$$n_r(s, b) = (p''_r(s, b))^2 \cdot \underbrace{(N(b) - n_c(s, b))}_{\text{bins available}} \qquad (10)$$

3.2.3. Total number of observed coincidences

The total number of observed coincidences then corresponds to the sum of the injected coincidences and the background coincidences:

$$n(s, b) = n_c(s, b) + n_r(s, b) \qquad (11)$$

For a comparison of simulations with the analytical description (Eq. (11)) see Fig. 4, solid lines. We note in

passing that, as long as the spike rates and coincidence rate are stationary across the observation time window, i.e. spikes and coincident spikes are uniformly distributed across time (which they are in this case), the mean and variance shown in Fig. 4 are independent of shifts of the binning window. Once the spike rates and/or coincidence rate vary as a function of time, and the spike trains (and/or coincidence trains) show a corresponding time structure, this may affect the values of mean and variance. The reason is that in such case there may be an interaction of the time structure in the spike trains and the positioning of the binning grid, depending on the various time constants involved. If the time constants of modulation in the spike/coincidence trains and the resolution of the binning grid become too similar, one may even approach a situation as observed in Moire-patterns or in aliasing. Hence, such choice of binning grid, effectively amounting to an undersampling of the relevant dynamics, should be avoided in the analysis.

3.3. Analytical description: expected number of coincidences

The predictor for the number of expected coincidences is based on the assumption of independence of the spike trains. Thus, it is based on the marginal probabilities. These are, however, a composition of the background firing probability and the coincident firing probability, reduced by the probability for clipped spikes due to the injection procedure (see Eq. (1)). The resulting firing probability p_{rc} is then modified by the binning (and clipping) procedure, as shown in Appendix C, Eq. (28) (see for illustration Fig. 5A, case of $p_c = 0$):

$$p_o = 1 - (1 - p_{rc})^b = 1 - (1 - (p_r + p_c - p_r \cdot p_c))^b \qquad (12)$$

The predictor for the expected number of coincidences corresponds to the product of the firing probabilities of the two trains, multiplied by the number of bins:

$$\hat{n} = \underbrace{(p_o)^2}_{\hat{P}} \cdot N(b) \qquad (13)$$

For small bin sizes, this function is dominated by the parabolic behavior of the expected joint probability (Fig. 5B). However for larger bin sizes, the decreasing number of bins $N(b)$ dominates (Fig. 5C), and reduces the number of expected coincidence counts considerably (Fig. 5D). Note, that this probability does not depend on the coincidence width in the data, it is a function of the bin size only. For comparison of simulations and the analytical description (Eq. 13) see Fig. 4, dotted line.

Coincidence Counts

Fig. 4. Coincidence analysis using the disjunct binning method (theory and simulation). Two parallel spike trains were generated with background rates $\lambda_r = 30$ sp/s, the rate of the injected coincidences was $\lambda_c = 5$ coinc/s, for $T = 10^6$ ms in time steps of $h = 1$ ms. The scatter of coincident events was varied from $s = \pm 0$ ms to $s = \pm 5$ ms. Each simulation experiment was repeated 10 times, and analyzed for the number of observed coincidences by varying the analysis bin width from $b = 1$ ms up to $b = 11$ ms. Results from repeated trials are indicated by the mean and variance (error bars). The solid lines indicate the analytical results. The dotted line shows the analytical result for the predicted number of coincidences, assuming independence of the two spike trains.

Fig. 5. Influence of disjunct binning on basic analysis parameters, here with $p_c = 0$. A: Spike probability p_o per bin (Eq. 12) as a function of the analysis bin size b, original spike probabilities p_r (per time step h) varied from 0.005 to 0.05. B: Expected probability (see $\hat{P}$ in Eq. 13) for coincidences as a function of bin size (original single spike probabilities as in A for both processes). C: Number of bins (Eq. 2) as a function of bin size; original time steps $N = 1000$. D: Number of expected coincidences (Eq. 13) as a function of bin size.
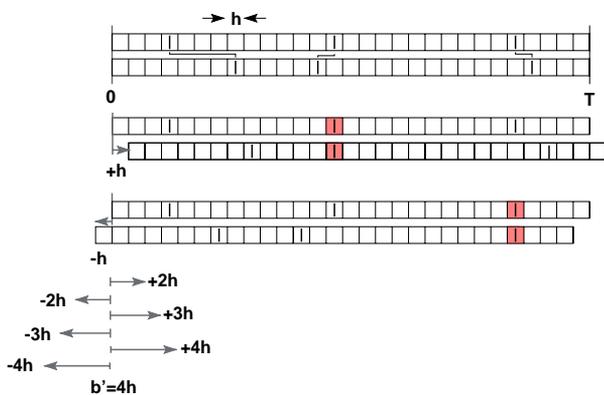
Fig. 6. Multiple Shifts. The second spike train is shifted against the first over up to $\pm b'$ in steps of the time resolution h (here shown for $b' = 4h$). For each shift, only exact coincidences (marked in grey) are counted. The sum over all shifts yields the total number of observed coincidences.

4. Multiple shifts (MS)

4.1. Detection of near-coincidences

In the alternative method of multiple shifts (cf. Fig.

6), the simultaneously recorded spike trains are analyzed for coincident events on their 'recording' time resolution h. Spikes that occur at the same time in the parallel spike trains are counted as coincident events. To account for near-coincidences, the second spike train is shifted against the first in steps of h up to $\pm b'$. For each shift, the 'exact' coincidences are counted, the sum over all shifts yields the observed coincidence count n.

For the predictor of the expected number of coincidences, the expected number of coincidences per shift of spike train 2 relative to spike train 1 is calculated and summed over all possible shifts $-b' \leq i \leq b'$:

$$\hat{n} = \sum_{i=-b'}^{b'} p_1 \cdot p_2(i) \cdot \frac{T}{h} \tag{14}$$

where $p_2(i)$ indicates that spike train 2 is shifted relative to spike train 1 by i steps on the original time resolution h. Since spike train 2 itself is not modified by shifting, the firing probability is unaffected, thus $p_2(i) = p_2$. In addition, we assume that the jitter of the near-coincidences is equally distributed over the total coincidence width. Then, the sum in Eq. (14) can simply be replaced by a factor of $2b' + 1$, and Eq. (14) can be rewritten as:

$$\hat{n} = p_1 \cdot p_2 \cdot \frac{T}{h} \cdot (2b' + 1) \tag{15}$$

For the significance evaluation of the difference between the observed and expected counts the joint-surprise (see Appendix B, Eq. 27) may be used again.

4.2. Analytical description: observed coincidences

4.2.1. Injected coincidences

All injected coincidences are detected, provided that the possible analysis shifts fully cover the scatter range $\pm s$ of the coincidences. At each single shift, only a fraction (on average: $1/(2s+1)$ of the injected coincidences is detected. Thus, per shift, the probability of detecting exact coincidences (p'_c) is given by this fraction of the total injection probability: $p_c \cdot 1/(2s+1)$. Thus, as long as the maximal analysis shift b' is smaller than the injected scatter width ($b' < s$), not all injected coincidences are detected, but only the fraction $(2b' + 1)/(2s + 1)$. For $b' = s$, the probability of detection is 1, for shifts larger than s, the probability is 0 (all coincidences are already detected for shifts up to $b' = s$). As a result, the coincidence probability becomes

$$p'_c(s, b') = \begin{cases} p_c \cdot \frac{2b' + 1}{2s + 1} & \text{for } b' < s \\ p_c & \text{for } b' \geq s \end{cases} \tag{16}$$

and the total number of detected injected coincidences is given by:

$$n_c(s, b') = p'_c(s, b') \cdot N \tag{17}$$

with $N = T/h$ the number of time steps of resolution h.

4.2.2. 'Background' coincidences

To obtain the background coincidences count, chance coincidences occurring at all possible time shifts (covering the analysis window $[-b', b']$) have to be summed. The effective background spike probability is composed of the original background activity p_r and possible 'debris' from not detected coincidences at a given analysis width (if $b' < s$). The contribution of debris per shift is $p_c \cdot (1 - 1/(2s + 1))$. For maximal analysis shifts $b' > s$, all spikes belonging to an injected coincidence are already counted as coincident spikes for shifts up to $b' = s$. Thus, for larger shifts, no debris are left as additional contributions to the background activity. In addition, a loss of background activity due to clipping by the injection procedure needs to be considered; this loss amounts to $p_r \cdot p_c$. As a result, the number of background coincidences is given by:

$$n_r(s, b')$$

$$= \begin{cases} \sum_{i=1}^{2b'+1} (p_r - p_r \cdot p_c + p_c \cdot (1 - \frac{1}{2s+1}))^2 \cdot N & \text{for } b' \le s \\ \sum_{i=1}^{2s+1} (p_r - p_r \cdot p_c + p_c \cdot (1 - \frac{1}{2s+1}))^2 \cdot N & \text{for } b' > s \\ \quad + \sum_{i=1}^{(2b+1)-(2s+1)} (p_r - p_r \cdot p_c)^2 \cdot N \end{cases}$$

$$= \begin{cases} (p_r - p_r \cdot p_c + p_c \cdot (1 - \frac{1}{2s+1}))^2 \cdot N \cdot (2b' + 1) & \text{for } b' \le s \\ (p_r - p_r \cdot p_c + p_c \cdot (1 - \frac{1}{2s+1}))^2 \cdot N \cdot (2b' + 1) & \text{for } b' > s \\ \quad + (p_r - p_r \cdot p_c)^2 \cdot N \cdot ((2b' + 1) - (2s + 1)) \end{cases} \tag{18}$$

4.2.3. Total number of observed coincidences

The total number of observed coincidences is the sum of the detected injected and the chance coincidences:

$$n(s, b') = n_c(s, b') + n_r(s, b') \tag{19}$$

For a comparison of simulations with the analytical description see Fig. 7, solid lines.

4.3. Analytical description: expected number of coincidences

As for DB, the firing probability (p_o) of a single spike train is the sum of the background activity (p_r) and the injected coincidence activity (p_c), reduced by the loss of spikes due to the injection procedure:

$$p_o = p_{rc} = p_r + p_c - p_r \cdot p_c \tag{20}$$

The predictor for the expected coincidence count is the product of the individual probabilities, summed over all possible time shifts and multiplied by the number of time steps:

$$\hat{n} = \sum_{i=1}^{2b'+1} (p_o)^2 \cdot N = \underbrace{(p_o)^2 \cdot (2b' + 1)}_{\hat{P}} \cdot N \tag{21}$$

The comparison of simulations with the analytical description for the expected number of coincidences is shown by the dotted line in Fig. 7.

Fig. 8 illustrates the behavior of the components that contribute to the expected number of coincidences. In contrast to the case of disjunct binning (see Fig. 5), the various components do not show a pronounced non-linear behavior. The firing probabilities are constant for increasing analysis shift sizes (Fig. 8A), the coincidence probability behaves, as expected, slightly parabolic (Fig. 8B), but does not go into saturation as for DB. Finally, the number of time steps is constant for increasing analysis shift size (Fig. 8C). Hence, the expected number of coincidences increases (slightly parabolic) (Fig. 8D), and does not show a maximum as in the case of DB.

5. Discussion

5.1. Comparison of methods: sensitivity for near-coincidences

The performance of the disjunct binning (DB) and the multiple shifts (MS) method is compared in relation

Fig. 7. Coincidence analysis using the multiple shifts methods (theory and simulation). Two parallel spike trains were generated with background rates $\lambda_r = 30$ sp/s and injected coincidence rate $\lambda_c = 5$ coinc/s, for $T = 10^6$ ms in time steps of $h = 1$ ms. The scatter of the coincident events was varied from $s = \pm 0$ ms to $s = \pm 5$ ms. Each simulation experiment was repeated 10 times, and analyzed for the number of coincidences by varying the maximal shift for analysis from $b' = \pm 0$ ms up to $b' = \pm 10$ ms. Mean and variance from repeated simulation experiments are shown by dots and error bars. Analytical results are shown by the curves (solid lines for the observed number of coincidences; dotted line: expected number of coincidences, assuming the two spike trains are independent).

Fig. 8. Influence of multiple shifts on basic analysis parameters, assuming $p_c = 0$. A: Spike probability p_o per time step as a function of analysis shift size b', original spike probabilities p_r varied from 0.005 to 0.05. B: Expected probability for coincidences ($\hat{P}$ in Eq. 21) as a function of b' (original spike probabilities as in A for both processes). C: Number of time steps as a function of b'; original time steps $N = 1000$. D: Number of expected coincidences (Eq. 21) as a function of b'. These results should be compared to Fig. 5 (disjunct binning).

to their ability to detect excess coincidences with respect to various coincidence widths of injected synchronous activity. The joint-surprise S is used as a measure for evaluating the sensitivity for excess coincidences. It compares the observed number of coincidences with the expected number of coincidences, based on the assumption of independent processes. Thus, the two methods are discussed in respect to these two components (n and $\hat{n}$), the difference between the observed and expected numbers of coincidences, and the resulting joint-surprise. Such a comparison using the same data parameters, i.e. background rate λ_r and coincidence rate λ_c, is shown in Fig. 9 for DB (A) and MS (B). The analyses using the analytical descriptions for DB and MS are illustrated for increasing analysis widths (bin size b for DB, maximal shift b' for MS). The corresponding analysis widths for DB and MS are $b' = b + 1$, since for a maximal coincidence spike distance s, the disjunct binning covers the distance within a bin of size $b = s + 1$. In the case for multiple shifts the distance s is covered by shifts up to $b' = s$.

Let us first consider the behavior of the joint-surprise S as a function of the analysis width. If coincidences were injected without any scatter ($s = 0$), both analysis

methods show a very similar behavior: the joint-surprise decreases non-linearly from its maximum at $b - 1 = 0$ for DB, and at $b' = 0$ for MS respectively. This decrease is due to the property of the joint-surprise, that evaluates the significance of excess coincidences in relation to the expected level. Thus, even for a constant level of excessive coincidences, S gets smaller with increasing b', because the expected level increases (see as an example Fig. 9B for $s = \pm 0$).

For coincidence widths $s > 0$, the S-functions for DB and MS differ. For small analysis widths ($b - 1 < s$, $b' < s$, respectively), the joint-surprise functions start for both, DB and MS, at low values smaller than the curve for $s = 0$, and they reach their maximum at $s = b - 1$ and $s = b'$, respectively. However, the joint-surprise curves are in general lower for DB, and do not converge to the same joint-surprise curve as they do in the case of MS. Consequently, its maxima are lower, and fall below the significance level of $\alpha = 0.01$ (line at $S = 2$ in Fig. 9, bottom graphs). As a consequence, the significance threshold is reached only for small values of the analysis bin width, i.e. at smaller analysis bin widths as compared to MS. For example, excess coincidences of a given coincidence width, e.g. $s = 4$ as in Fig. 10, are only detected as significant events by MS, but not with DB.

The reason for these differences are best visualized by the differences of the observed and the expected number of coincidences (Fig. 9, third row from top). For $b - 1 < s$ (or $b' < s$) only a fraction of the injected coincidences are detected in both methods, although for different reasons: in case of DB, since coincidences are split and therefore cannot be detected, in case of MS, only a fraction of the coincidences are covered by the time shift. However, from $b' = s$ on, all coincidences are detected in case of MS. In the example, the difference $n - \hat{n}$ reaches a constant value at 93.3986, a value slightly smaller than the number of injected coincidences (here: 100). This is a result of a negative difference between the (accumulated) background coincidences (up to the maximal shift) (Eq. 18) and the predicted number of coincidences (Eq. 21). In contrast, in case of DB the differences between the observed and the expected number of coincidences are in general lower than for MS. They even decrease for large bin sizes, which is a result of binning followed by clipping and the strong reduction of the number of bins (see Fig. 4C,D).

Summarizing, the splitting of coincidences by disjunct binning at small bin sizes, combined with the strong non-linear effect of binning followed by clipping for larger bin sizes leads to a considerably lower sensitivity for excess coincident events as compared to the MS. Thus, multiple shifts are better suited for the detection of near-coincidences for any coincidence width.
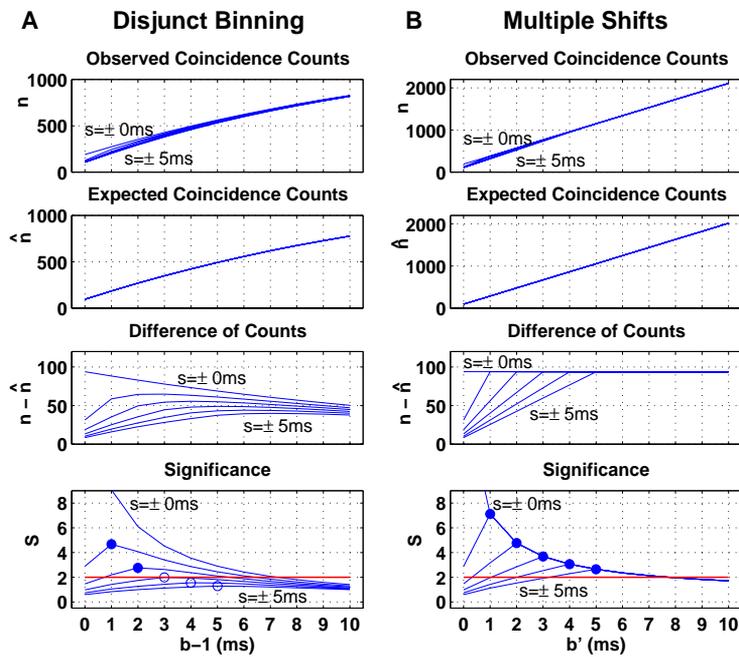
Fig. 9. Comparison of the performance of disjunct binning (A) and multiple shifts (B) using analytical descriptions. For various coincidence widths ($s = \pm 0 \ldots \pm 5$ ms) the observed coincidence counts (top row), the expected coincidence counts (second row), the differences of the two (third row) and the joint-surprise (bottom row) are calculated for increasing analysis widths (in A: for bin sizes b from 1 to 11, in B: for maximal shifts b' from 0 to 10). For better comparison, in A and B the abscissae are represented such that they correspond to the same maximal spike distances of a coincidence, that is covered by the analysis widths, i.e. $b = b' + 1$. Parameters were $\lambda_r = 30$ sp/s, $\lambda_c = 1$ coinc/s, $T = 10^5$ ms, $h = 1$ ms. The maxima of the joint-surprise curves are marked by filled circles, if the values are above the significance level of $\alpha = 0.01$, i.e. $S = 2$; if they are below they are marked by empty circles. The maxima of S for $s = \pm 0$ (not shown) are at $b - 1 = 0$, $S = 16.76$ in A and at $b' = 0$, $S = 16.73$ in B.

5.2. The temporal precision of neuronal processes

In order to detect the temporal precision in an example set of real neuronal data, we will now use the multiple shifts method. In a first step, experimental data are analyzed for the significance of coincident spike events (by using the joint-surprise) for various analyses shift widths b'. In a second step, the coincidence firing probability p_c of the experimental data is calculated from the model, introduced in Section 2 for comparison of the two analysis methods. In a third step, experimental results are compared to simulated data, in order to test the validity of our model assumptions.

Neuronal data were taken from a pair of simultaneously recorded neurons (time resolution: $h = 1$ ms) from the primary motor cortex of a behaving monkey involved in a visuo-manual pointing task (see Bastian et al. (1998) and Grammont and Riehle (1999) for experimental details). A first stimulus instructed the monkey about the required movement direction, a second stimulus after a delay of 1000 ms called for its performance. From repetitive trials (here: 33) of the same behavioral condition, a time segment of 800 ms during which the rates of the neurons were stationary, was cut beginning 100 ms after the first stimulus and analyzed for coincident spike events.

The results of the MS analyses of the experimental data using various shift widths b' are shown in Fig. 11 (solid lines). The number of observed coincidences (Fig. 11A) and the number of expected coincidences (Fig. 11B) increase approximately linear with b'. The joint-surprise function, however, shows a clear peak at $b' = 6$ ms. In the left column of Fig. 11, the experimental results are compared to control surrogate data (shown in grey), in which no coincidences were injected. The
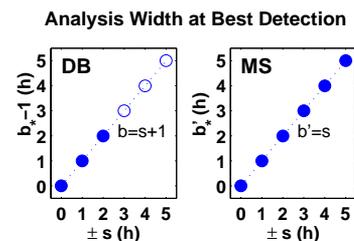
Fig. 10. Detectability of near-coincidences: comparison of disjunct binning (left) and multiple shifts (right). The analysis width $b_* - 1$, b'_* at the maximum of the joint-surprise S, as extracted from Fig. 9, is plotted against the injected scatter width $\pm s$. Its linear relation for both analysis types (DB and MS) shows that the analysis width at the maximum of the joint-surprise indeed indicates the underlying coincidence width s of the data. However, in case of DB, for larger analysis bin sizes (here: $b_* - 1 > 2$) coincidences are not detected as significant anymore (empty circles), in contrast MS does (filled circles).
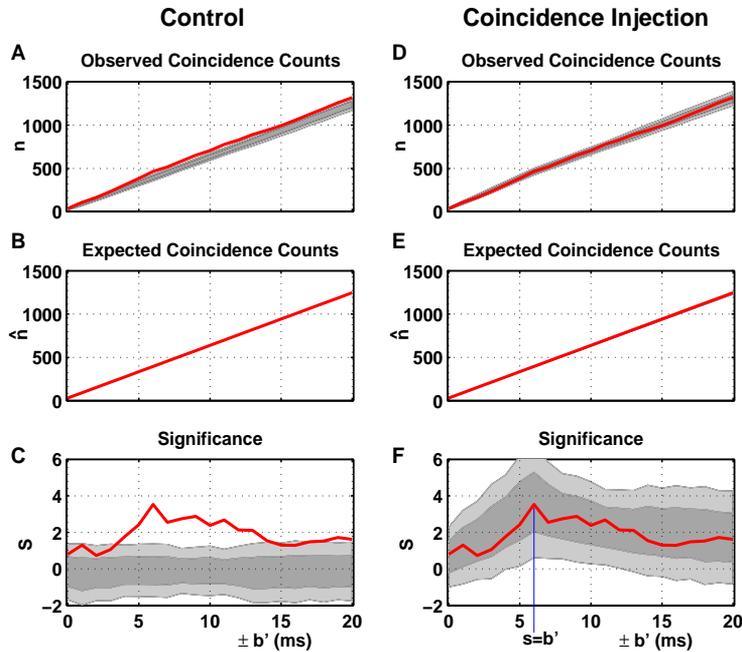
Fig. 11. Temporal precision of neuronal spike trains and its comparison with simulated processes. In both columns, the analysis results for increasing shift width b' of two simultaneously recorded neurons are shown and compared to simulated data. Top row: observed number of coincidences n, middle row: expected number of coincidences $\hat{n}$, bottom row: joint-surprise S. In the left column, for control purposes the simulation experiment is performed without injected coincidences, firing probabilities correspond to the measured marginals of the neurons ($p_1 = 0.0321$, $p_2 = 0.0359$). On the right, coincidences were injected with a coincidence width of $s = 6$ ms, corresponding to analysis shift width b' at the maximum of the joint-surprise. The coincidence and background probabilities were calculated based on the model shown in Fig. 1 ($p_c = 0.0029$, $p_r^1 = 0.0291$ and $p_r^2 = 0.0329$). Results from simulations are shown as grey bands. The width of the light grey band represents 95%, the dark grey band 70% of 30 repetitions of the simulation experiments. Each simulation had the same duration in time as the experimental data (33 trials of 800 ms) with a time resolution of $h = 1$ ms. In panel C the light grey band (representing 95% of the simulation experiments) is well below the threshold for significance of 1%, demonstrating the low probability of the significance measure to generate 'false alarms'.

rates for the simulations were set to correspond to the marginal firing rates of the neuronal data. Simulation results of 30 repetitions (each consisting of 33 trials) are shown as grey bands. To visualize the density of the resulting distribution, the range of 95% is shown in light grey, the range of 70% in dark grey. The experimental results clearly deviate from the simulation results, indicating that neuronal spike trains do not correspond to the assumption of independence. The light grey band in Fig. 11C (representing 95% of the simulation experiments) is well below a conservative threshold for significance of 1%, demonstrating the low probability of the significance measure to generate 'false alarms' (i.e. high specificity).

As discussed above, according to our model the maximum of S in the experimental data indicates the coincidence width of the underlying data as $s = b'_*$, i.e. here 6 ms. Next, surrogate data with injected coincidences were compared to the experimental results. According to our results shown in Fig. 9, we extracted the coincidence width for the simulation at the maximum of the joint-surprise, i.e. $s = 6$ ms. The firing probabilities of the neurons, measured as the marginal probabilities, were assumed to be a measure for the sum of

coincident and background activity. Thus, for different rates of the two neurons i we obtain

$$p_i = p_c + p_r^i, \quad \text{for } i = 1, 2 \tag{22}$$

From the observed number of coincidences at $s = 6$ ms, the coincidence probability p_c can be calculated by expressing Eq. (19) using Eq. (17) and Eq. (18) as follows. The number of coincident spikes is given by Eq. (17) as $n_c(s, s) = p_c \cdot (2s + 1)/(2s + 1) \cdot N$. Chance coincidences are derived using Eq. (18) (but here for $p_r^1 \neq p_r^2$). Thus, we restate Eq. (19) as:

$$\begin{aligned}
&n(s, s) \\
&= p_c \cdot \frac{2s + 1}{2s + 1} \cdot N + \left(p_r^1 - p_r^1 \cdot p_c + p_c \cdot \left(1 - \frac{1}{2s + 1} \right) \right) \\
&\quad \cdot \left(p_r^2 - p_r^2 \cdot p_c + p_c \cdot \left(1 - \frac{1}{2s + 1} \right) \right) \cdot N \cdot (2s + 1)
\end{aligned} \tag{23}$$

By rearranging Eq. (23) and replacing p_r^i by use of Eq. (22) we obtain a function of 4th order in p_c. Neglecting terms of 3rd and 4th order, we get as an expression for p_c the positive root:

$$p_c = -\frac{B}{2A} + \sqrt{\left(\frac{B}{2A}\right)^2 - \frac{C}{A}} \qquad (24)$$

with

$$A = 1 + 2 \cdot (p_1 + p_2) + p_1 \cdot p_2$$
$$- \left(1 - \frac{1}{2s+1}\right) \cdot (p_1 + p_2 + 2) + \left(1 - \frac{1}{2s+1}\right)^2$$

$$B = \frac{1}{2s+1} - p_1 - p_2 - 2 \cdot p_1 \cdot p_2$$
$$+ \left(1 - \frac{1}{2s+1}\right) \cdot (p_1 + p_2)$$

$$C = p_1 \cdot p_2 - \frac{n}{(2s+1) \cdot N} \qquad (25)$$

Using the measured marginal probabilites p_1 and p_2, we obtain estimates for the coincidence probability $p_c = 0.0029$ and the background firing probabilities $p_r^1 = 0.0291$ and $p_r^2 = 0.0329$ (per time step h).

Fig. 11D–F illustrates the comparison of experimental and simulated data using the above derived parameters. The experimentally derived joint-surprise function shows basically the same curve as obtained from the surrogate data. The experimental curve lies within the range of 70% (about 1σ) of the simulated data (dark grey band in Fig. 11F), indicating that our model predictions are consistent with the experimental data.

Because the multiple shifts method recovers more of the coincidences actually hidden in the data than disjunct binning does, it necessarily increases the sensitivity of the UE-method. Even though Fig. 11F is not primarily concerned with sensitivity — rather, it compares physiological data with surrogate data generated by a specific model — relevant conclusions regarding the sensitivity can be drawn from it. Note, that not all cases with injected coincidences are detected as UEs: there is a certain percentage of false negatives. To visualize this more clearly, we indicated the distribution of the surrogate data by different grey levels (light grey indicating the 95% range, dark grey the 70% range). This visualization of the distribution once more underlines the importance of choosing an adequate time resolution for the analysis: when adopting a conservative threshold of 0.01, about 15% of the simulations with injected coincidences are not detected as UEs at an optimal analysis resolution, for non-optimal time resolution the amount of false negatives may even be considerably higher. Compared to the specificity of the method (cf. Fig. 11C), the sensitivity (Fig. 11F) is distinctly lower at a threshold level of 0.01. Obviously, on the basis of model simulations as in Fig. 11C and F, a more optimal choice of threshold level can be tuned that simultaneously optimizes sensitivity and specificity (Grün, 1996; Grün and Aertsen, 1999a).

6. Conclusions

We examined two alternative methods, the disjunct binning and the multiple shifts, to detect excessive coincidences and their coincidence width. The performance of these methods was tested using surrogate data sets. In the case of disjunct binning, the number of detected coincidences is reduced considerably (for small bin sizes) by splitting coincidences due to the application of the binning grid. This effect decreases with increasing bin size. Binning followed by clipping leads to an increase of occupied bins. This reduces the number of observed coincidences as well as the amount of expected coincidences, and yields a non-linear increase of both counts for increasing bin sizes. For analysis bin sizes larger than the injected coincidence width ($b - 1 > s$), the difference of the observed and the expected number of coincidences decreases non-linearly with increasing bin size, and thus leads to a reduced sensitivity for excess coincident events. These two effects are antagonistic and lead to a maximum of the significance at $b - 1 = s$.

In contrast, in the case of the multiple shifts, the number of expected coincidences increases linearly with increasing analysis bin width. For analysis bin widths smaller than the injected scatter ($b' < s$), a fraction of the coincidences cannot be detected. However, if the maximal analysis shift equals the coincidence width ($b' = s$) all injected coincidences are detected. For larger shifts ($b' > s$) also chance coincidences (due to background activity) reach a considerable count. At $b' = s$ the difference of actual and expected coincidences reaches its maximum, leading to a maximum of the joint-surprise function. From there on ($b' > s$) the difference remains constant, since both measures increase further (linearly) by the same amount of background coincidences. This leads to a decrease of the significance (note significance curve for $\pm s = 0$, for $b' \geq s$ all curves are identical).

Both methods allow to detect the temporal scatter of near-coincident spike events in neuronal data. They detect near-coincidences best at analysis widths just covering the width of the injected scatter (i.e. disjunct binning at $b = s + 1$, multiple shifts at $b' = s$). Both methods are comparable for $b' = b + 1$. However, for higher temporal scatter only the multiple shifts method is suited to detect excess coincidences reliably as significant events. Thus, the multiple shifts is a tool to detect the precision of coincident activity in neuronal processes, and, in addition, may be used to derive the underlying coincidence rate.

Appendix A. List of symbols used

T	temporal duration of the observation interval, $[T] =$ unit of time
h	original time resolution, $[h] =$ unit of time
N	total number of bins
b	size of a bin in the disjunct binning method in units of h, $[b] = 1$
b'	largest absolute time shift in the multiple shift method in units of h, $[b'] = 1$
s	largest absolute time difference between spikes in an injected coincidence in units of h, $[s] = 1$
λ_r, λ_c	background, coincidence rate, $[\lambda_r] = [\lambda_c] = 1/$unit of time
p_s	probability for the 2nd coincidence spike to be placed at any of the positions within the scatter range
p_r, p_r^j	probability of finding a bin occupied, contributions of background, rate for neuron j
p_c	probability of finding a bin occupied, contributions of coincidence rate
p_c'	probability for finding a coincidence in MS
p_{rc}	probability of finding a bin occupied for the combined processes (background and injected)
p_o	analysis method dependent marginal probability for the combined processes
$\hat{P}$	expected joint probability
p_j	marginal firing probability of neuron j
n, n_r, n_c	number of coincidences, contributions of background and injected coincidences
$\hat{n}$	expected number of coincidences
F	fission probability for near-coincidences in the disjunct binning method
p_r'	probability of finding a bin (on the original time scale) occupied by a back ground spike, corrected for debris and collision for DB
p_r''	probability of finding a bin occupied by a background spike after binning and clipping (corrected for debris and collision) for DB
S	joint surprise function
b_*, b_*'	analysis width at which the joint-surprise is maximal
P	cumulative probability of observing the measured number of coincidences or even a larger one
α	significance level
p_b	probability for finding a bin occupied after binning and clipping

Appendix B. The 'Unitary Events' method (Grün, 1996)

Action potentials (spikes) elicited by M simultaneously recorded neurons are transformed, using an appropriate binning, to M-dimensional joint-activity vectors consisting of ones (spike) and zeros (no spike). The composition of these vectors represents the various constellations of coincident spiking activity across the M neurons (measured coincidences). Under the null-hypothesis that the M neurons fire independently, the expected number of occurrences of any joint spike constellation is calculated as the product of the individual firing probabilities. The measured number of coincidences (n) is then compared with the expected number of coincidences ($\hat{n}$).

The statistical significance for a positive difference between the numbers of measured and expected coincidences is calculated from a Poisson distribution (with the mean being set to the expected coincidence number) as the cumulative probability P of observing the measured number of coincidences (or an even larger one) by chance.

$$P(n|\hat{n}) = \sum_{r=n}^{\infty} \frac{\hat{n}^r}{r!} \cdot e^{-\hat{n}} \qquad (26)$$

The larger the number of excess coincidences, i.e. more coincidences are measured than expected by chance, the closer P approaches 0. Similarly, the larger the number of lacking coincidences, i.e. less coincidences are measured than expected by chance, the closer its complement $1 - P$ approaches 0, while P approaches 1. In order to enhance visual resolution at the relevant low probability values of P or $1 - P$, we calculate a logarithmic function, the 'joint-surprise' (S), of the two:

$$S(P) = \log_{10} \frac{1-P}{P} \qquad (27)$$

It is derived from the 'surprise' measure (Palm, 1981; Palm et al., 1988; Aertsen et al., 1989) and is comparable to measuring significance on a dB-scale. For excessive coincidences, this function is dominated by P, for lacking coincidences by $1 - P$. One obtains positive values for excessive coincidences, negative values for lacking coincidences, and zero at chance level ($P = 0.5$). Whenever the significance value of an excess number of measured coincidences exceeds a fixed threshold α (here: $\alpha = 0.01$, corresponding to a joint-surprise value of 2), this defines an epoch containing significantly more coincidences than expected by chance. These instances of excessive coincident spiking are referred to as epochs with 'Unitary Events' (UEs).

The dependence of the sensitivity/specificity of the UE method on the firing rates and the coincidences contained

in the data is discussed in Grün (1996) and Grün et al. (1999a).

In order to account for non-stationarities in the discharge rates of the observed neurons, modulations in spike rates and coincidence rates are determined on the basis of short data segments by sliding a fixed time window (typically 100 ms wide) along the data in steps of the coincidence binwidth. This timing segmentation is applied to each trial, and the data of corresponding segments in all trials are then analyzed as one quasi-stationary data set, using the appropriate estimation of the time-varying rates. Conventional approaches to estimate time-varying firing rates are based on trial-averaging (PSTH; Gerstein and Kiang, 1960). This, however, makes the strong assumption of stationarity across trials. If this condition is not fulfilled, more sophisticated analyses are required to estimate the dynamic firing rates from single-trial data, on which to base the expected coincidence counts. We are currently evaluating the applicability of a newly developed method for single-trial rate estimation (Nawrot et al., 1999) for this purpose.

Appendix C. Spike reduction by binning and clipping

In order to be able to treat a binned spike train as a binary process, we clip the data to 1 in the case of more than one spike within a bin. In a single bin, the probability of finding any event combination (i.e. 0's and 1's) is given by the binomial distribution, summing up to 1. Thus, the probability of finding one or more spikes (k) within a single bin $b \cdot h$, given the probability p for a spike per time step h on the original spike train, is given by:

$$p_b(k \geq 1) = \sum_{k=1}^{b} \binom{b}{k} p^k \cdot (1-p)^{b-k} = 1 - (1-p)^b \quad (28)$$

As a result, the marginal probability of finding a 1 in a bin is increased as compared to the original spike train, converging to 1 for very large bin sizes (see Fig. 5A).

References

Abeles M. Local Cortical Circuits: An Electrophysiological Study. Berlin: Springer-Verlag, 1982.

Abeles M. Corticonics: Neural Circuits of the Cerebral Cortex, 1st edition. Cambridge: Cambridge University Press, 1991.

Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Thishby N, Vaadia E. Cortical activity flips among quasi stationary states. Proc Nat Acad Sci USA 1995;92:8616–20.

Abeles M, Bergman H, Margalit E, Vaadia E. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. J Neurophysiol 1993;70:1629–38.

Aertsen A, Gerstein G, Habib M, Palm G. Dynamics of neuronal firing correlation: Modulation of 'effective connectivity'. J Neurophysiol 1989;61:900–17.

Aertsen A, Vaadia E, Abeles M, Ahissar E, Bergman H, Karmon B, Lavner Y, Margalit E, Nelken I, Rotter S. Neural interactions in the frontal cortex of a behaving monkey: signs of dependence on stimulus context and behavioral state. J Hirnf 1991;32:735–43.

Bastian A, Riehle A, Erlhagen W, Schöner G. Prior information preshapes the population representation of movement direction in motor cortex. NeuroReport 1998;9:315–9.

Gat I, Tishby N, Abeles M. Hidden markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. Network Comp Neural Sys 1997;8:297–322.

Gerstein GL, Aertsen AMHJ. Representation of cooperative firing activity among simultaneously recorded neurons. J Neurophysiol 1985;54:1513–28.

Gerstein GL, Bedenbaugh P, Aertsen A. Neuronal assemblies. IEEE Trans Biomed Eng 1989;36:4–14.

Gerstein GL, Kiang NYS. An approach to the quantitative analysis of electrophysiological data from single neurons. Biophys J 1960;1:15–28.

Grammont F, Riehle A. Precise spike synchronization in monkey motor cortex involved in preparation for movement. Exp Brain Res 1999; 128:118–22.

Grün S. Unitary Joint-Events in Multiple-Neuron Spiking Activity: Detection, Significance, and Interpretation. Reihe Physik, Band 60. Thun, Frankfurt/Main: Verlag Harri Deutsch, 1996.

Grün S, Diesmann M, Aertsen A. 'Unitary Events' in multiple-single-neuron activity. I. Detection and significance, 1999a. (to be submitted).

Grün S, Diesmann M, Aertsen A. 'Unitary Events' in multiple-single-neuron activity. II. Non-stationary data, 1999b. (to be submitted).

Hebb DO. Organization of behavior. A neurophysiological theory. New York: Wiley, 1949.

Nawrot M, Aertsen A, Rotter S. Single-trial estimation of neuronal firing rates. J Neurosci Meth 1999;94:81–91.

Palm G. Evidence, information and surprise. Biol Cybern 1981;42:57–68.

Palm G. Cell assemblies as a guideline for brain research. Conc Neurosci 1990;1:133–48.

Palm G, Aertsen A, Gerstein GL. On the significance of correlations among neuronal spike trains. Biol Cybern 1988;59:1–11.

Papoulis A. Probability, Random Variables, and Stochastic Processes, 3rd edition. New York: McGraw-Hill, Inc, 1991.

Riehle A, Grün S, Diesmann M, Aertsen A. Spike synchronization and rate modulation differentially involved in motor cortical function. Science 1997;278:1950–3.

Singer W. Synchronization of cortical activity and its putative role in information processing and learning. Annu Rev Physiol 1993;55:349–74.

von der Malsburg C. The Correlation Theory of Brain Function. Internal report 81–2. Göttingen, FRG: Max-Planck-Institute for Biophysical Chemistry, 1981.