Journal of Neuroscience Methods 179 (2009) 90–100

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

Unbiased estimation of precise temporal correlations between spike trains

Eran Stark[a,*], Moshe Abeles[a,b,c]

[a] Department of Physiology, Hadassah Medical School, Hebrew University, Jerusalem 91120, Israel
[b] Interdisciplinary Center for Neural Computation, Hebrew University, Jerusalem 91904, Israel
[c] Gonda Brain Research Center, Bar Ilan University, Ramat-Gan 52900, Israel

ARTICLE INFO

Article history:
Received 14 October 2008
Received in revised form
26 December 2008
Accepted 29 December 2008

Keywords:
Correlation analysis
Extra-cellular recordings
Jittering
Neuronal interactions
Poisson distribution
Resampling methods
Spike trains
Temporal precision

ABSTRACT

A key issue in systems neuroscience is the contribution of precise temporal inter-neuronal interactions to information processing in the brain, and the main analytical tool used for studying pair-wise interactions is the cross-correlation histogram (CCH). Although simple to generate, a CCH is influenced by multiple factors in addition to precise temporal correlations between two spike trains, thus complicating its interpretation. A Monte-Carlo-based technique, the jittering method, has been suggested to isolate the contribution of precise temporal interactions to neural information processing. Here, we show that jittering spike trains is equivalent to convolving the CCH derived from the original trains with a finite window and using a Poisson distribution to estimate probabilities. Both procedures over-fit the original spike trains and therefore the resulting statistical tests are biased and have low power. We devise an alternative method, based on convolving the CCH with a partially hollowed window, and illustrate its utility using artificial and real spike trains. The modified convolution method is unbiased, has high power, and is computationally fast. We recommend caution in the use of the jittering method and in the interpretation of results based on it, and suggest using the modified convolution method for detecting precise temporal correlations between spike trains.

1. Introduction

To reveal mechanisms underlying the operation of neuronal networks *in vivo*, neurophysiologists typically seek to understand the activity of individual neurons and inter-neuronal interactions. A key analytical tool used to study inter-neuronal interactions is the cross-correlation histogram (CCH; Perkel et al., 1967). For a pair of spiking neurons, the CCH counts the number of times the two neurons spiked at exactly the same instant and at various times one relative to the other. Similar tools have been developed for analyzing the concurrent activity of more than two spiking neurons (Perkel et al., 1975; Abeles, 1983; Prut et al., 1998; Abeles and Gat, 2001; Schrader et al., 2008).

Despite the extensive use of the CCH, its interpretation is not trivial because a non-flat CCH may result from precise temporal interactions between two neurons but also from slower temporal co-variability (Brody, 1999) and/or from correlations with stimulus or action parameters (Ben-Shaul et al., 2001). To reliably interpret precise temporal features in multi-neuronal activity and in CCHs, a statistical manipulation of the spike trains has been suggested according to which the precise time of each spike is jittered within a small window (Date et al., 1998). Analyses based on jittering have recently been used to uncover novel operating mechanisms of cortical networks (Hatsopoulos et al., 2003; Shmiel et al., 2005; Fujisawa et al., 2008).

Upon using the jittering method we found that it is lacking in the sense that statistical tests based on it are conservative and have a low sensitivity. Here, we describe these findings, give a mathematical rationale for them, and suggest an alternative method which is unbiased, has high sensitivity, and is computationally fast. We demonstrate the method using artificial spike trains and real spike trains of single-units (SUs) recorded from the Macaque premotor cortex.

2. Materials and methods

2.1. Cross-correlation methods

2.1.1. Computation of a raw CCH

Given two spike trains each quantified as a 0/1 series in time, the cross-correlation histogram measures the number of times spikes from the two trains occurred at the same instant and at various time lags relative to one another. If the bin width of the histogram is B (s), then the CCH count at any bin m in the range $\pm M$ is the number of spikes from one (trigger) train followed, within $mB - B/2$ to $mB + B/2$, by a spike from the other (reference) train.

* Corresponding author. Current address: Center for Molecular and Behavioral Neuroscience, Rutgers University, 197 University Avenue, Newark, NJ 07102, United States.
E-mail address: eranst@andromeda.rutgers.edu (E. Stark).

0165-0270/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jneumeth.2008.12.029

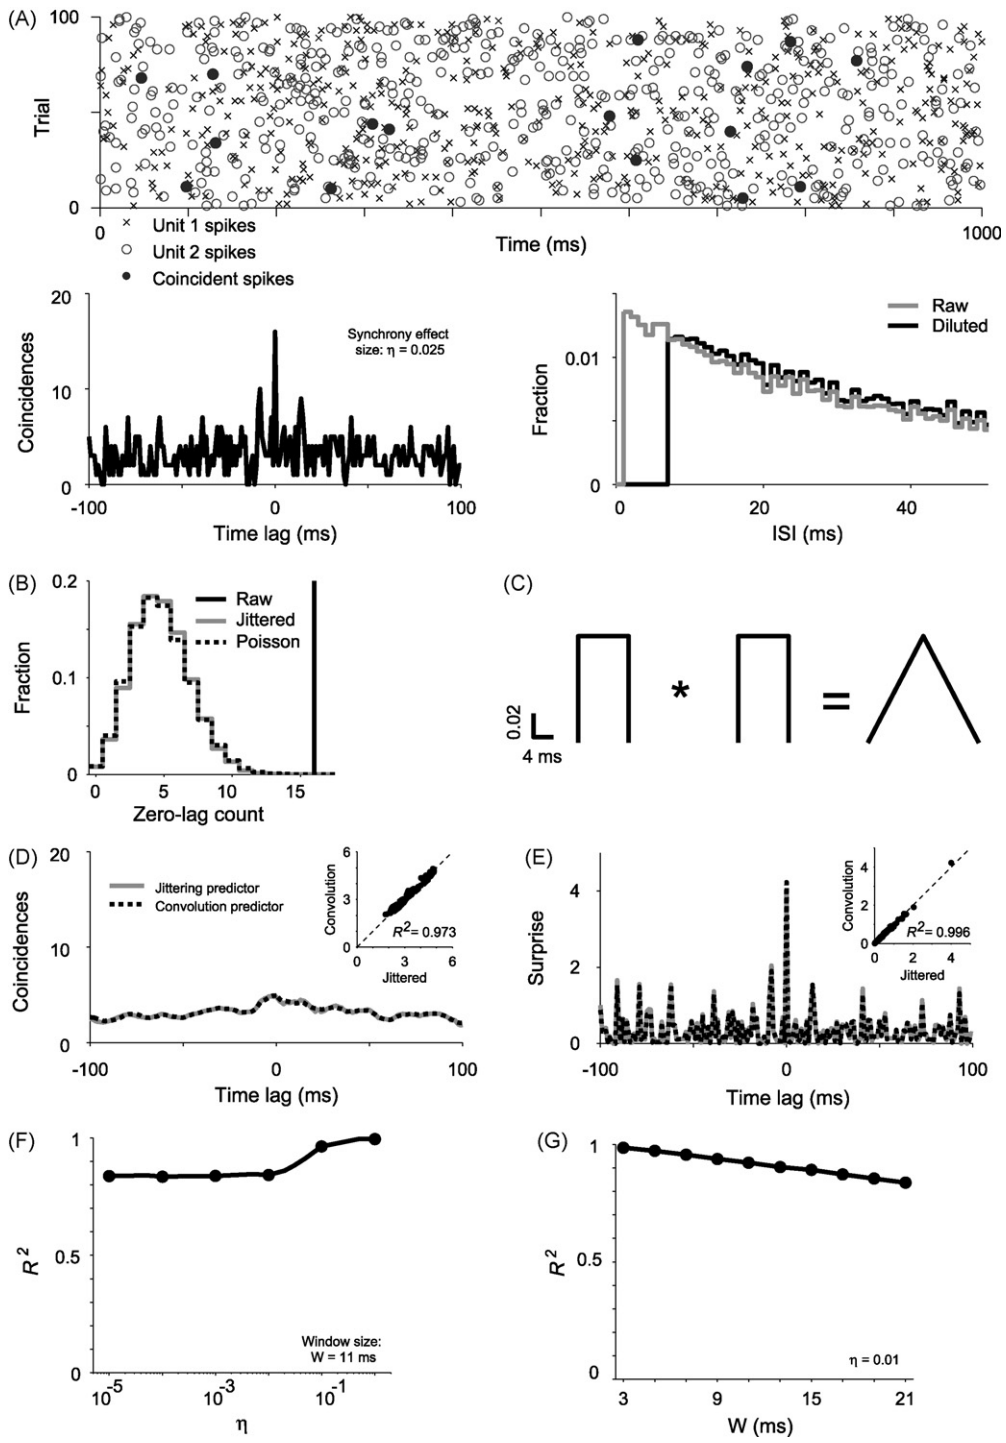

Fig. 1. Jittering is asymptotically identical to a Poisson assumption. (A) Simulated data. *Top*: spike trains. Each of the two putative neurons fires at a rate of 5 spikes/s; 100 trials each 1 s long were simulated, with a total of 488 (unit 1) and 519 (unit 2) spikes. The two spike trains are slowly correlated (see Eq. (2)) and synchronized; synchrony effect size is $\eta = 0.025$, indicating that ~2.5% of the spikes occur at the same millisecond (Eq. (4)). *Bottom left*: raw CCH for the simulated trains. *Bottom right*: inter-spike interval (ISI) histogram for the simulated data (unit 1). Following dilution, the minimum ISI is 7 ms. (B) Distribution of zero-lag counts. Counts are shown for the raw CCH (black line, same as in the zero-lag bin of (A), *bottom left*); for 10,000 CCHs derived from jittered spike trains (gray histogram); and for the Poisson distribution with parameter $\lambda = 4.78$ spikes/s (dashed black histogram), determined by convolving the raw CCH with the triangular window shown in (C). (C) Rectangular window used for jittering individual spikes in each of the two trains. The convolution of two rectangular windows of width W (here, $W = 11$ ms) yields a triangular window of width $2W - 1$ used for smoothing the raw CCH. (D) Predictor CCHs, obtained by jittering both spike trains 10,000 times and averaging (continuous gray line) or by convolving the raw CCH with the triangular window (dashed black line). *Inset*: scatter plot of the bin-by-bin predictor values, showing that the two predictors are indistinguishable. (E) p-Value CCHs, shown as Surprise $= -\log_{10}$ p values (Palm et al., 1988). The p-value at a given bin tests the hypothesis that the count in that specific bin of the raw CCH is above the corresponding value in the predictor CCH. Other conventions are the same as in (D). *Inset*: scatter plot of the bin-by-bin surprise values. The two methods yield essentially identical results, although the jittering method has lower resolution, limited by the number of jittering repetitions (see the deviating point comparing the p-values of the zero-lag bin). (F) Correspondence between the jittering and convolution methods as a function of η, varied between 10^{-5} (nearly uncorrelated spike trains) and 1 (identical trains). For each value of η, 1000 spike train pairs were simulated as in (A). The R^2 measures the bin-by-bin match of the two predictors. Regardless of η the R^2 is high, indicating that the two methods yield nearly identical results. (G) Correspondence between the jittering and convolution methods as a function of W, the jittering window width. $\eta = 0.01$; other conventions are the same as in (F). When the window size is small relative to the minimum ISI, the two methods yield very similar results.

For instance, when $B = 0.001$ s and $M = 100$, the zero-lag ($m = 0$) bin includes spikes that occurred "simultaneously" at time lags of −0.5 to 0.5 ms, and the last ($m = 100$) bin includes spike pairs where a reference spike occurred 99.5–100.5 ms after the trigger spike. This yields the raw CCH, a vector of $2M + 1$ non-negative integers that count the number of (possibly lagged) coincident spikes of two spike trains (e.g. Fig. 1A, *bottom left*).

There are two delicate points in this procedure. First, to keep the counting process in each bin Poissonian even when the spike trains themselves are not, spike trains should be diluted prior to CCH computation (Abeles and Gat, 2001). Here, this was done by deleting all spikes except the first one whenever an inter-spike interval (ISI) was below $R = 6$ ms. Second, for a finite measuring duration, the raw CCH computed as explained above has an asymptotically triangular shape for uncorrelated trains and is therefore particularly biased to high values in the zero-lag bin. To see why, consider a pair of spike trains each $T = 1$ s long; the count in the zero-lag bin is based on 1000 samples, whereas counts in flanking bins are based on less and less samples: the highest lagging bins are based on only $T/B - M = 900$ samples. To prevent this counting bias we use only $T/B - M$ samples for all lags. This is done by selecting one train as a trigger, keeping only the first $T/B - M$ samples from it, and computing the CCH with the entire reference train (call this CCH_1). Then the process is repeated with the roles of the two trains switched (CCH_2). The right-hand sides of the two CCHs are then combined: the $1, 2, \ldots, M$ bins from CCH_2 are reversed in time before concatenating with the $0, 1, \ldots, M$ bins of CCH_1, so that bin numbers $1, 2, \ldots, M$ of CCH_2 are bin numbers $-1, -2, \ldots, -M$ of the compound (unbiased by sample size) CCH.

2.1.2. Null hypothesis for precise temporal correlations and statistical testing

The raw CCH simply counts the number of (possibly lagged) coincident spikes in two trains. To interpret a CCH, we employ the null hypothesis H_0 "the two spike trains are uncorrelated at a precise temporal resolution". The key word here is "precise"; for instance, for a temporal resolution of 5 ms, CCH features 50 ms wide are not considered a rejection of H_0. A main objective of the current work is to develop an unbiased and sensitive procedure for testing this null hypothesis. Although the following applies to any bin in the histogram, for simplicity we will focus specifically on the zero-lag bin of the CCH which measures synchrony, the number of times spikes from the two trains occurred simultaneously up to a temporal resolution B.

Several methods have been proposed to test the abovementioned null hypothesis (Perkel et al., 1967; Abeles, 1982b; Aertsen et al., 1989; Baker et al., 2001), with the following common steps. First, compute a raw CCH. Second, compute a predictor CCH. Third, based on some probability distribution, estimate the probability to obtain the raw CCH (or some features of it) given the predictor CCH.

For a set of identically and independently distributed (i.i.d.) trials, a predictor CCH can be obtained by shuffling (or shifting) the order of one set of spike trains (trials) relative to the other (Perkel et al., 1967; Aertsen et al., 1989). Since there are many possible permutations, shuffling is typically repeated several times and the resulting CCHs are averaged, yielding the "shuffle" predictor, a vector λ of $2M + 1$ non-negative scalars that count the number of (possibly lagged) coincident spikes of two spike trains assuming that spike trains are i.i.d. The probability to obtain an observed (or higher) count n in the mth bin (time lag) of the raw CCH, given the value $\lambda(m)$ in the same bin of the predictor CCH, is estimated using the Poisson distribution (Abeles, 1982b):

$$P(n \text{ or more}|\lambda(m)) = 1 - \sum_{x=0}^{n-1} \left(\frac{e^{-\lambda(m)}\lambda(m)^x}{x!} \right). \quad (1)$$

However, experimentally collected data are rarely i.i.d. because there are often various uncontrolled sources of trial-to-trial variability (Brody, 1999; Ben-Shaul et al., 2001). Thus, the trial-shuffling procedure is unsuitable for distinguishing between precise temporal correlations and correlations that occur on a slower (intra- or inter-trial) temporal resolution. Moreover, sometimes the experimental paradigm does not consist of a set of well-defined trials of equal lengths and conditions but rather the data form a continuously collected stream with no "trials" that may be shuffled. Therefore, trial-shuffling is inadequate for testing our null hypothesis.

2.1.3. The jittering method

An alternative approach that directly tests our H_0 is to randomly vary the time of each spike within a small "jittering" window (Date et al., 1998; Abeles and Gat, 2001; Hatsopoulos et al., 2003). The width W of the jittering window determines the temporal precision: for instance, when all spikes from both trains are jittered using a uniform (rectangular) window of W (ms), "precise temporal correlation" is any CCH feature spanning less than $2W$ (ms). A predictor CCH conforming to H_0, without precise temporal correlations, is then obtained by jittering, computing a jittered CCH, repeating this many times, and averaging the jittered CCHs. The sample of jittered CCHs also provides an empirical non-parametric distribution on which H_0 can be evaluated: under the null hypothesis, the probability to obtain an observed or higher count in the zero-lag bin of the raw CCH is equal to the fraction of jittered CCHs with the observed or higher count in the zero-lag bin.

When jittering spike trains, a couple of technical points must be considered. First, either one or both spike trains can be jittered (Pazienti et al., 2007). Since our null hypothesis is that spike timing does not matter whatsoever up to some temporal resolution, both trains are jittered within the same uniform (rectangular) window with $W = 11$ ms unless stated otherwise, implying a temporal precision of ±5 ms. Second, it has been suggested that spike trains should be jittered with respect of the inter-spike intervals (ISIs; Gerstein, 2004). When ISIs are small relative to W, jittering may cause two consecutive spikes to (1) become closer than a refractory period of R (ms), (2) overlap, hence decimating one, or even (3) switch their chronological order. For a rectangular window, decimation may occur when $ISI < W$ and the first spike is jittered forward in time and the second backwards. Similarly, a jittered ISI smaller than R samples may occur when $ISI < W + R$. To reduce the occurrence of such cases, W should be small relative to the minimum ISI; to completely prevent violation of refractoriness for any R and W, the jittering procedure is modified as follows. Whenever the ISIs to the previous and the next spike are above $W + R$, the spike is jittered within a window of width W. When intervals are shorter, the shortest stretch of consecutive spikes with flanking intervals above $W + R$ is detected and all spikes within this stretch are jittered repeatedly until no violation of refractoriness R is generated.

2.1.4. The convolution method

A second method that directly tests our H_0 is the convolution method, developed for the detection of precise firing patterns (Abeles and Gat, 2001) and adapted here to the study of precise pair-wise correlations as follows. First, compute a raw CCH. Second, compute a predictor CCH by convolving the raw CCH with a window of width W, essentially smoothing out sharp CCH features. To prevent edge effects, the first $W/2$ bins (excluding the very first bin) are duplicated, reversed in time, and prepended to the CCH prior to convolving. Likewise, the last $W/2$ bins are symmetrically appended to the CCH. As in the jittering method, W determines the temporal precision of the null and alternative hypotheses. Third, using Eq. (1), estimate the probability to obtain an observed or higher count in the raw CCH given the predictor CCH.

Table 1
Summary of default parameter values.

Usage	Parameter	Notation	Value	Units
Simulations	Dilution order	R	6	ms
	Mean firing rate	F	5	spikes/s
	Noise S.D.	σ	0.5	spikes/s
	Number of trials	N	100	–
	Temporal resolution	Δt	0.1	ms
	Time constant	τ	50	ms
	Trial length	T	1	s
CCH computations	Bin width	B	0.001	s
	Maximal lag	M	100	bins
	Rectangular window width	W	11	ms

2.2. Data source

The data used in this work are of two types: artificial and real spike trains. We used artificial spike trains because then spiking statistics such as firing rates, number of trials, slow temporal correlations, temporal drifts, and synchrony effect sizes can be accurately controlled. We used spike trains recorded from the Macaque premotor cortex because then no assumptions regarding the statistics of the individual trains have to be made.

2.2.1. Simulated spike trains

Spike train pairs were simulated without any correlations (H_0), with slow temporal correlations (still H_0), and with precise temporal correlations (with and without slow temporal correlations; H_1). Unless stated otherwise, simulations proceeded in time steps of $\Delta t = 0.1$ ms and $N = 100$ i.i.d. trials each lasting $T = 1$ s were generated. For a summary of the parameters used in this study, see Table 1. Results were insensitive to specific parameter values.

First, a firing rate profile was generated. For independent spike trains, a fixed profile of $f(t) = F_{\text{ind}}$ spikes/s was used. For slowly correlated trains, we employed a time-varying profile:

$$f(t) = e^{-1/\tau} f(t - \Delta t) + \text{Noise}(t) \tag{2}$$

with initial conditions $f(0) = 0$ and zero-mean Gaussian noise, $\text{Noise}(t) \sim N(0,\sigma)$ (we used a time constant of $\tau = 50$ ms and a noise S.D. of $\sigma = 0.5$ spikes/s). This produced a distinct slowly changing rate profile for each trial with zero mean and auto-correlation that falls exponentially with τ. The rate profile was translated to have a mean F_{ind} and then negative values were clipped to zero.

Second, given an underlying firing rate profile $f(t)$, a Poisson spike train was simulated by deciding at each time step whether a spike did or did not occur. This was done by random sampling from a binomial distribution with parameters $B(n,p) = (1, f(t)\Delta t)$. In other simulations, spike trains with non-Poisson ISI distributions (gamma processes of order k) were generated by decimating the Poisson spike trains, keeping only every kth spike (Baker and Gerstein, 2000). The latter process was repeated twice, resulting in a pair of independently generated (uncorrelated or slowly correlated) spike trains for each trial.

Third, to simulate precise temporal synchrony of a predetermined magnitude, identical (synchronous) spikes were inserted into the two independently generated trains (Grün et al., 1999). The "common source" spike train was generated using a fixed rate profile F_{common} and subsequently merged with each of the two independently generated trains. Assuming that the final desired mean firing rate of each spike train is F, the mean rates of the common and independent processes are given by

$$\begin{aligned} F_{\text{common}} &= \eta F \\ F_{\text{ind}} &= (1 - \eta) F, \end{aligned} \tag{3}$$

where the synchrony effect size $0 \le \eta \le 1$ is defined as

$$\eta = \frac{\text{Number of coincident spikes}}{\text{Total number of spikes}}. \tag{4}$$

The effect size η does not distinguish between slow and precise temporal correlations, depends on the CCH bin width B, and may differ for the two spike trains if the total number of spikes differs between the two trains.

2.2.2. Cortical spike trains

Real spike trains were recorded from the premotor cortices of two Macaque monkeys that participated in a prehension task. The task, recording procedures, and signal processing techniques are described elsewhere (Stark et al., 2007, 2008) and are only briefly summarized below. All animal handling procedures were in accordance with the *NIH Guide for the Care and Use of Laboratory Animals* (1996), complied with Israeli law, approved by the Ethics Committee of the Hebrew University, and supervised by a veterinarian. During each recording session, 6 reach directions and 2 grasp types were used for a total of 12 distinct reach-grasp combinations. After the monkey pressed a button, one object was shown at one location for 200–400 ms, indicating the upcoming reach direction and grasp type. After a delay of 1000–1500 ms during which the monkey was required to continue pressing the button and no visual stimulus was provided, a go signal prompted the monkey to reach, grasp, and hold the target object (for at least 580 ms) without visual feedback. Correct trials were reinforced by a juice reward. Trial length, defined as the time from cue onset to the end of the hold period, was 2675 ms (median; 95% range: 2348–3048).

Premotor neural activity was recorded using up to 16 independently movable glass-coated tungsten microelectrodes (impedance: 0.2–2 MΩ at 1 kHz) arranged in two circular guide tubes (8 electrodes in each; Double MT, Alpha-Omega Engineering, Nazareth, Israel). During each session, one tube was aimed towards arm-related (shoulder or elbow) regions of the dorsal premotor cortex (PMd) and another towards finger-related regions of the ventral premotor cortex (PMv). From the raw electrical signal recorded by each electrode (amplified 10K, band-pass filtered 1–10,000 Hz, and sampled at 25 kHz), spikes were detected using a modified second derivative algorithm and sorted into well-isolated single-units (SUs; Stark et al., 2007). SUs recorded by different electrodes were analyzed if the mean firing rate over the 2500 ms following cue onset was >0.5 spikes/s and the unit was recorded for at least 30 trials per reach-grasp combination exhibiting stationary activity (Stark et al., 2008). A total of 112 SUs conformed to these criteria, forming 629 jointly recorded spike train pairs.

2.3. Chance level CCH count and synchrony effect size

For two uncorrelated spike trains, the expected count λ at the zero-lag bin (or at any other bin) depends only on spiking probability and recording duration. For two trains with stationary firing

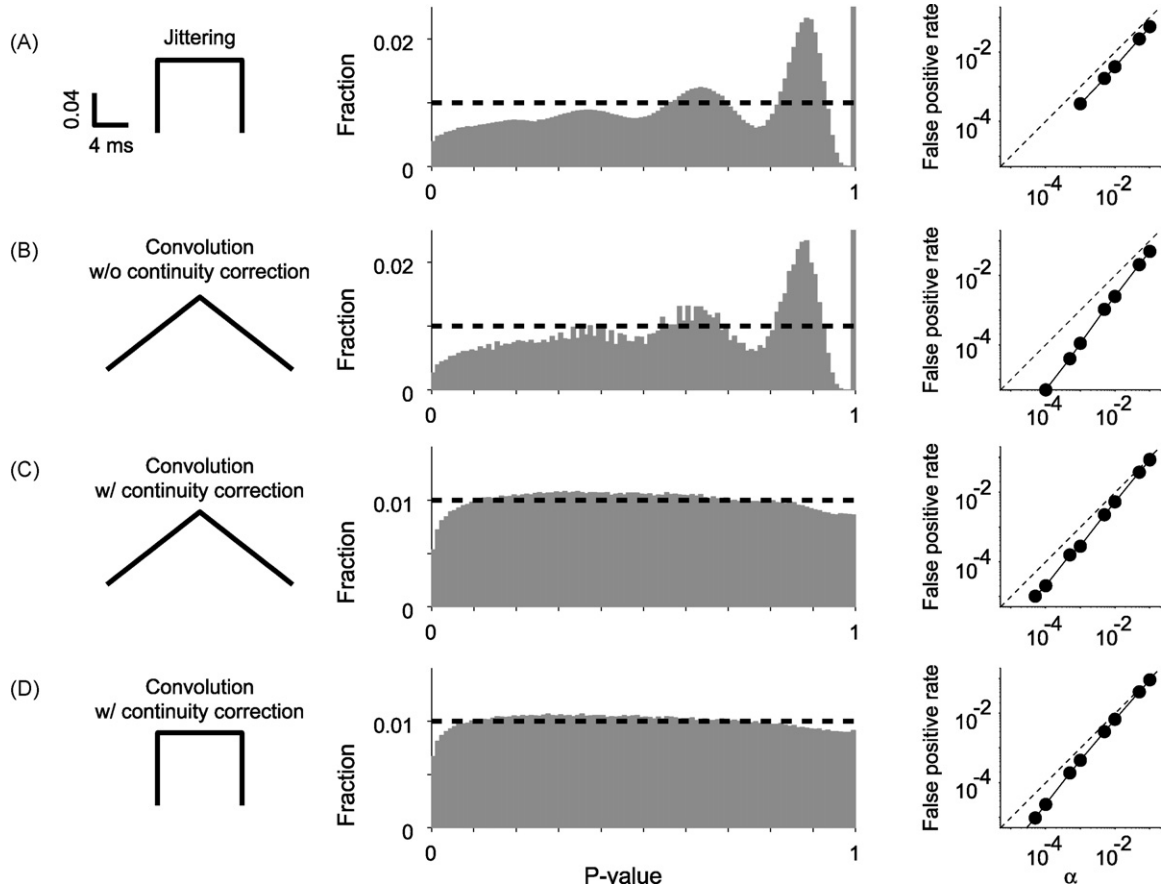

Fig. 2. Jittering and convolution are conservative tests. (A) *Left*: jittering window. *Center*: *p*-Value distribution of all CCH bins (as in Fig. 1D) in a sample of 1000 pairs of uncorrelated spike trains. For each pair, *p*-values were determined by jittering spikes 1000 times. Abscissa shows *p*-values, ordinate shows the fraction of values within each bin, and the dashed horizontal line shows the expected fraction in each bin (bin size is 0.01). The distribution is highly non-uniform: there is a paucity of low *p*-values, and 12% of the *p*-values are in the right-most bin (between 0.99 and 1; the ordinate was clipped for presentation purposes). *Right*: false positive rate, the fraction of correlated CCH bins derived from uncorrelated spike trains. The empirical false positive rate (plotted on a log-scaled ordinate) is consistently lower than the α level (plotted on a log-scaled abscissa), indicating that the jittering method is conservative. (B) *p*-Values determined using the straight-forward convolution method (triangular window). Data and conventions are the same as in (A). The jittering and convolution methods yield similar *p*-value distributions and both are conservative. (C) *p*-Values determined using the convolution method (triangular window) with a continuity correction (see Section 3). The distribution is more uniform than without a correction (B), but the number of extreme (low and high) *p*-values is below the expected. (D) *p*-Values determined using the convolution method with a continuity correction using a rectangular window (equivalent to jittering only one of the spike trains). As for a triangular window, the test is overly conservative.

rates F_1 and F_2 spikes/s, the expected value is

$$\lambda = F_1 F_2 T_E B, \tag{5}$$

where T_E is the effective recording duration and B is the CCH bin width (both measured in s). The exact expression is slightly more complex because of the procedure used to correct the CCH for finite recording duration (see Section 2.1.1). For N trials each T (s) long, T_E is shorter than $N \times T$; for a maximal CCH lag of M samples, $T_E = N(T - BM)$. Thus $\lambda = F_1 F_2 N(T - BM)\ B$. For instance, using the same $F = 4.85$ spikes/s for both trains (Poisson trains with 5 spikes/s, with ISIs up to $R = 6$ ms diluted) and other parameters as in Table 1 we obtain $\lambda = 2.12$.

Synchrony effect size η can be estimated even without the generative assumption of common and independent processes and thus this measure can be employed for quantifying synchrony in pairs of real spike trains. However, even for uncorrelated and unsynchronized trains, η is not necessarily zero. Assuming two independently and stationary spike trains with an identical mean firing rate F, Eq. (5) can be rewritten as $\lambda = F^2 T_E B$. Plugging the latter expression into Eq. (4) yields $\eta_{\text{chance}} = (\lambda / F T_E) = (F^2 T_E B / F T_E) = FB$. When the firing rates of the two trains are different, chance-level η can be estimated by the geometric mean:

$$\eta_{\text{chance}} = B\sqrt{F_1 F_2}. \tag{6}$$

For the parameters in Table 1 η_{chance} is 0.005.

3. Results

3.1. Jittering is asymptotically equivalent to CCH smoothing and a Poisson assumption but is a conservative test

The zero-lag bin of the raw cross-correlation histogram counts the number of times that spikes from two trains occurred within the same time bin; for an example of weakly synchronized artificial spike trains and the corresponding CCH, see Fig. 1A. Jittering each spike within a small rectangular window enables testing the null hypothesis that the two trains are uncorrelated up to a well-defined temporal resolution. For the example CCH, the raw count at zero-lag (black line in Fig. 1B) was higher than all jittered values (10,000 jitters; gray line in Fig. 1B). Thus, the null hypothesis of unsynchronized spike trains was rejected in this case with a *p*-value smaller than 1/10,000.

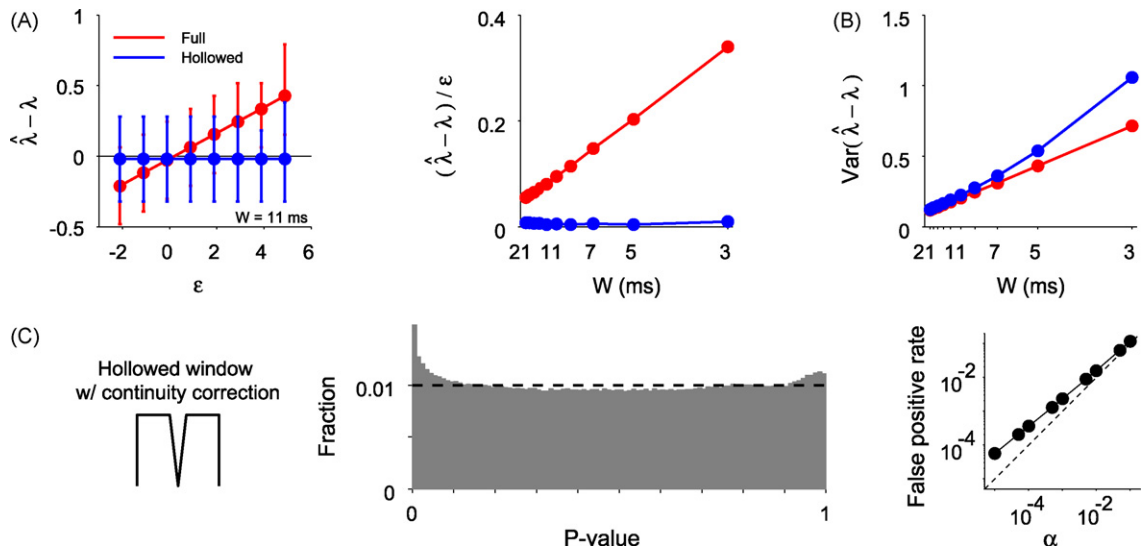

Fig. 3. Convolution with a hollowed window is a permissive test. (A) Bias in the estimation of the expected value $\hat{\lambda}$ of the zero-lag CCH bin. 20,000 CCHs were generated based on uncorrelated spike trains (as in Fig. 2). $\hat{\lambda}$ was estimated for each CCH by convolving the raw CCH with full and hollowed rectangular windows. *Left*: the median bias $\hat{\lambda} - \lambda$ in the estimate, plotted as a function of the deviation ε of the observed count C from λ ($\varepsilon = C - \lambda$; Section 3.2). Error bars show lower and upper quartiles; abscissa clipped at $\varepsilon = 6$ for presentation purposes. Convolution with a full window yields a biased estimate of λ with a slope of $1/W$ (shown here for $W = 11$ ms), whereas a hollowed window yields unbiased estimates. *Right*: the slope $(\hat{\lambda} - \lambda)/\varepsilon$, plotted for various window sizes. Sample is the same as in (A). The hollowed window yields unbiased estimates of λ regardless of W. (B) Variance in the estimation of $\hat{\lambda}$. The variance was computed using full and hollowed windows of various widths. For each W, the hollowed window yields more variable estimates than the full window. Sample is the same as in (A). (C) p-Values determined using the convolution method employing a hollowed rectangular window. Sample and conventions are the same as in Fig. 2. The fraction of low (and high) p-values is above the expected, indicating that the test based on a hollowed window is permissive.

Because convolution and correlation are linear operators, their order is interchangeable. When the minimal inter-spike interval is sufficiently large (see below), jittering each of two spike trains within a rectangular window of width W is asymptotically equivalent to convolving the raw CCH with a triangular window of width $2W - 1$ (Fig. 1C). Therefore, the jitter CCH predictor (based on jittering each spike within a rectangular window) and the convolution CCH predictor (based on convolving the raw CCH with the appropriate triangular window) are essentially the same ($R^2 = 0.97$; F-test: $p \ll 0.001$; Fig. 1D). Moreover, because the distribution of the jittered counts is essentially Poisson (Kolmogorov–Smirnov test for the zero-lag bin: $p = 0.19$; compare gray and dashed lines in Fig. 1B), significance tests based on the two methods yield nearly identical results. Specifically, the bin-by-bin probability of above-chance coincident counts is the almost the same for the jittering and convolution methods ($R^2 > 0.99$; F-test: $p \ll 0.001$; Fig. 1E). The jittering method is computer intensive and the resolution of the resulting p-value depends on the number of jittering repetitions, whereas the convolution method is computationally fast and yields high-resolution p-value (compare the p-value for the zero-lag bin; deviating point in Fig. 1E, *inset*). The two methods yield near-identical predictors (and p-values; data not shown) regardless of the synchrony effect size η (Fig. 1F). However, the correspondence between the predictors is reduced for spike trains with small ISIs relative to the jittering/smoothing window width (Fig. 1G). Thus, for testing the null hypothesis that "the two spike trains are uncorrelated at a precise temporal resolution", convolution of the raw CCH yields equivalent results to jittering the original spike trains, provided that the minimum ISI is large relative to the convolution window width.

Jittering is a conservative test in the sense that for uncorrelated spike trains, the p-value distributions are highly non-uniform and show three notable features (Fig. 2A, *center*): an excess of values close to 1; a complex pattern of deficient/excess intermediate values; and a deficit of small p-values. This latter feature means that the false positive rate is lower than the significance threshold α for each α used (Fig. 2A, *right*). Since jittering with a rectangular window and smoothing with a triangular window are asymptotically equivalent (Fig. 1), the latter procedure has the same features (Fig. 2B). The main reason for the first two non-uniform features is the discontinuous nature of the CCH count (here, at zero-lag): because proper probability estimation requires an integral transformation of a continuous random variable, estimating the probability of a discrete random variable is prone to bias (Pearson, 1950). It was therefore suggested to add a small random number to the discrete variable, thereby obtaining a continuous random variable (Pearson, 1950). For an observed CCH count n, the continuity-corrected p-value is a random number between the Poisson probability to see (n or more) and ($n + 1$ or more) counts (Abeles and Gat, 2001). This procedure removes the very high p-values and the complex pattern of intermediate values, resulting in a flatter p-value distribution. However, there is still an underestimate of small (and large) p-value (Fig. 2C). The same result is obtained when a rectangular window is used for smoothing the CCH (Fig. 2D). This means that the test based on convolution and a continuity correction is still conservative.

Intuitively, the jittering and convolution methods are both biased even after applying a continuity correction because the corresponding predictors over-fit the spike data (in the sense that the observed number of coincidences at a given time lag is used to estimate chance-level information at the same lag). More formally, to see why jittering results in over-fitting, consider two uncorrelated spike trains with an expected count λ in all CCH bins. For simplicity, assume that we estimate the convolution predictor for the zero-lag bin using a rectangular window of width W (as in Fig. 2D). In practice, the observed count C in that bin deviates from λ by a random number ε; because C distributes as a Poisson process with parameter λ, $<\varepsilon> = 0$ and $<\varepsilon^2> = \lambda$ independently of W. The total count under a window centered on the zero-lag bin is then $C + (W-1)\lambda = (\lambda + \varepsilon) + (W-1)\lambda = W\lambda + \varepsilon$, and the estimated predictor value at zero-lag is $\hat{\lambda} = (W\lambda + \varepsilon)/W = \lambda + \varepsilon/W$. Thus, although asymptotically converging to λ, $\hat{\lambda}$ is always biased in the direction of ε so the estimated expected value deviates in the same direction as the observed count (Fig. 3A). In other words, the probability estimation of the observed count in a given CCH bin is

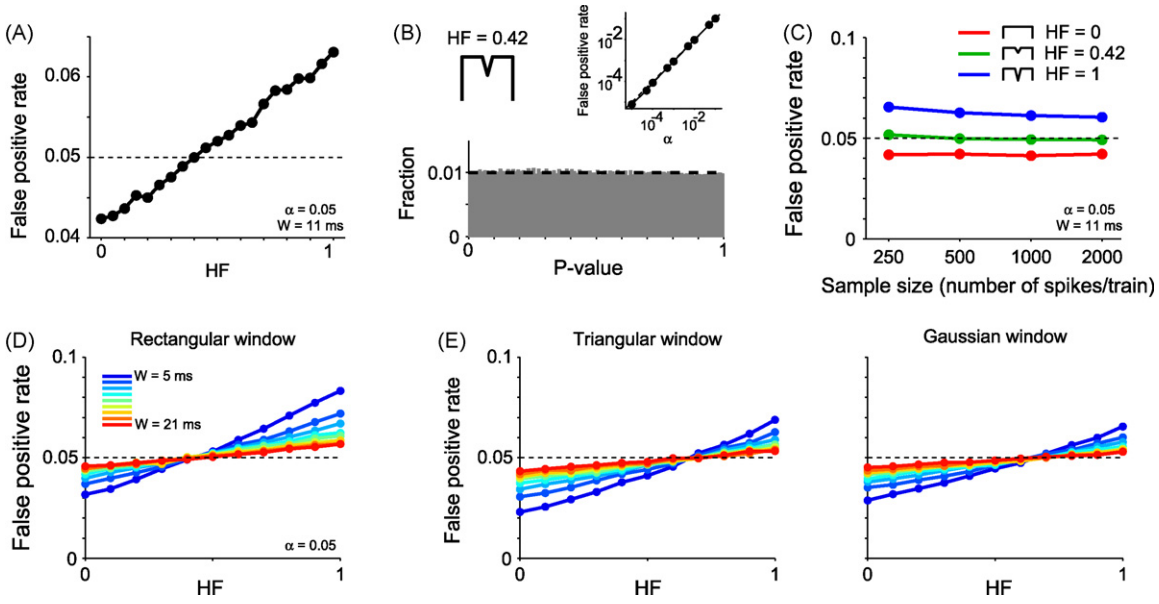

Fig. 4. Convolution with a partially hollowed window is unbiased. (A) False positive rate vs. hollowed fraction. Continuity-corrected *p*-values were estimated using a rectangular window ($W = 11$ ms) with hollowed fractions (HFs) between 0 (full window) and 1 (hollowed window). Sample is the same as in Fig. 2. The expected rate for uncorrelated data and an unbiased test is α. The HF corresponding to a false positive rate of 0.05 is 0.42. (B) *p*-Values determined using the modified convolution method, with a partially hollowed rectangular window (HF = 0.42). Sample and conventions are the same as in Fig. 2. The *p*-value distribution is uniform and the false positive rate is identical to the α level for each α tested, indicating that the test is unbiased. (C) False positive rate as a function of sample size. For a partially hollowed window, the false positive rate does not depend on the number of spikes used in the analysis. (D) Optimal HF for rectangular windows of various widths *W*. Although the optimal HF does not depend on *W*, the false positive rate is especially sensitive to the HF at small window sizes. (E) Optimal HF for triangular and Gaussian windows. Conventions are the same as in (D). Regardless of *W*, the optimal HF is 0.6 for a Gaussian window and 0.63 for a triangular window.

biased (Fig. 2) because it is based on a biased predictor that over-fits the spike trains.

3.2. Convolution with a hollowed window is a permissive test

The bias in the estimate of the expected CCH count results from over-fitting the original spike data (Section 3.1), and therefore a possible way to eliminate this bias is to minimize the contribution of the value in the bin in question (here, the zero-lag bin in the CCH) to the estimation of the expected value in that bin (Abeles and Gat, 2001). Consider a rectangular window of width *W* with the central bin of the window set to zero. For the same setting as before (uncorrelated spike trains with expected value λ in all CCH bins and a zero-lag bin with $\lambda + \varepsilon$ counts), the total count under the hollowed window is $(W-1)\lambda$ and the predictor value is $\hat{\lambda} = [(W-1)\lambda]/(W-1) = \lambda$. Thus, in contrast to a full window, the estimate of the expected value in a CCH based on a hollowed rectangular window is independent of ε and hence unbiased (Fig. 3A, *left*) regardless of window width *W* (Fig. 3A, *right*).

However, because the error in the estimation of the expected value (the variance of $\hat{\lambda}$) is higher when a hollowed window is employed (Fig. 3B), probability estimation based on a hollowed window yields an excess of small and large *p*-value. To see this, consider that for uncorrelated spike trains the estimate $\hat{\lambda}_{\text{full}}$ (based on a full rectangular window of width *W*) is the mean value of *W* random variables, each distributing Poisson with parameter λ. Thus, the expected value of the estimator is $E(\hat{\lambda}_{\text{full}}) = \lambda$ and its variance is $Var(\hat{\lambda}_{full}) = \lambda/W$. While the estimate based on a hollowed window has the same expected value $E(\hat{\lambda}_{\text{hollowed}}) = \lambda$, its variance is larger, $Var(\hat{\lambda}_{\text{hollowed}}) = \lambda/(W-1)$. This effect is especially apparent for small *W*s (Fig. 3B). Thus although $\hat{\lambda}_{\text{hollowed}}$ is uncorrelated with the specific deviation ε, it can have more extreme values than $\hat{\lambda}_{\text{full}}$ and therefore yields more extreme *p*-values when used for estimating probabilities (Fig. 3C, *center*). Consequently the empirical false positive rate of the convolution method based on a hollowed window is higher than the test α (for each α level tested; Fig. 3C, *right*), indicating that the test is permissive.

3.3. Convolution with a partially hollowed window yields an unbiased test

Because full and hollowed windows lead to conservative and permissive tests, respectively (Sections 3.1 and 3.2), we decided to employ a partially hollowed window as a potential compromise. We investigated the possibility that a partially hollowed window will yield an unbiased test by systematically varying the hollowed fraction (HF) of a rectangular window between 0 (full window) and 1 (hollowed window). For each HF we computed the false positive rate for CCHs derived from uncorrelated spike trains (Fig. 4A). We found that a partially hollowed rectangular window with HF = 0.42 (Fig. 4B, *left*) yields a false positive rate identical to the test α regardless of α (Fig. 4B, *right*). Specifically, the *p*-value distribution for uncorrelated data using such a window is flat (Fig. 4B, *bottom*).

The partially hollowed window yielded unbiased *p*-value distributions regardless of sample size (the number of spikes in the analyzed spike trains; Fig. 4C). While the optimal HF was 0.42 independently of the window size *W*, the dependence of the false positive rate on the HF was especially pronounced for narrow windows (small *W*; Fig. 4D). Although qualitatively identical results were obtained for other window types (Fig. 4E), the optimal HFs differed, being 0.6 for Gaussian windows and 0.63 for triangular windows. This is consistent with the fact that these windows, when not hollowed, give a higher weight to the value in the central bin (the triangular more than the Gaussian) and thus over-fit the raw CCH to a higher extent than a rectangular window.

An important issue in statistical testing is the test power, the complement of the false negative rate $(1-\beta)$. Power is essentially the hit rate, the probability to determine that an effect is present when it actually is; the higher the power, the more sensitive a test is. In general, power depends on both the effect size η and the significance threshold α, decreasing when α is low (Neyman and

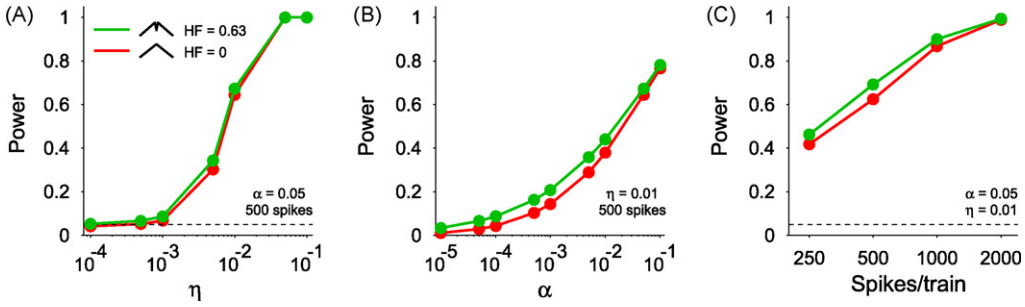

Fig. 5. Power of the modified convolution method. (A) Test power vs. synchrony effect size η. For each η, 10,000 spike train pairs were simulated and the power (hit rate) was computed for the zero-lag (synchronized) bin. The test based on a partially hollowed triangular window has higher power than the test based on a full window, increasing the sensitivity to small effects. (B) Test power vs. α level. η is 0.01; other conventions are the same as in (A). Regardless of α, the test based on a partially hollowed window has higher power than the test based on a full window. (C) Test power vs. sample size. Regardless of the number of spikes per train, the test based on a partially hollowed window has higher power.

Pearson, 1933). Indeed, the power of the convolution method based on a triangular window with a continuity correction increases with η (Fig. 5A), α (Fig. 5B), sample size (Fig. 5C), and the HF (Fig. 5). For weakly synchronized ($\eta = 0.01$) artificial neurons for which ~2000 spikes were generated (each neuron firing at 5 spikes/s and measured over 400 one-second long trials), the modified convolution method enables detection in $96.5 \pm 0.6\%$ (mean $\pm$ S.E.) of the cases when $\alpha = 0.01$. Using $\alpha = 0.05$ increases the power, enabling detection in $99.3 \pm 0.3\%$ of the cases (Fig. 5C). Thus, under these circumstances the test misses in less than 1% of the cases.

In sum, the modified convolution method, using a partially hollowed window and a continuity correction, yields a false positive rate which is equal to the α level regardless of α and is thus unbiased. This is true for various window types and stands in contrast to the other techniques considered here including the jittering and the convolution methods with/without a continuity correction (Fig. 6A). Finally, the power of the modified convolution method based on a partially hollowed window is high relative to the other methods, with the highest power obtained when a Gaussian window is employed (Fig. 6B).

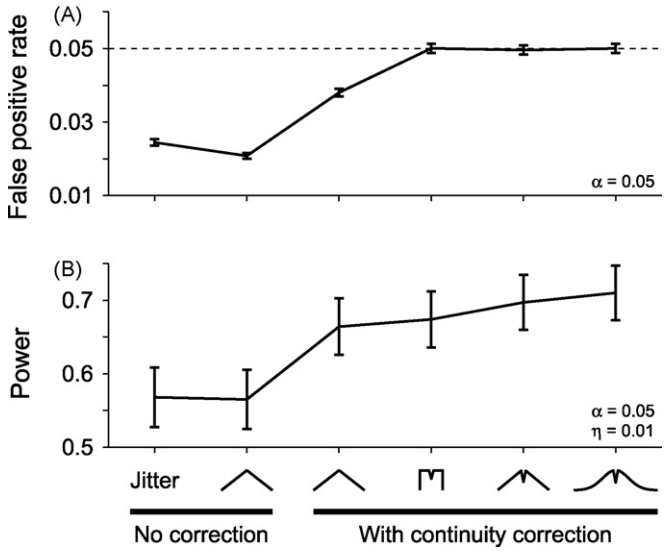

Fig. 6. Summary of the comparison between the different methods. (A) False positive rate of selected tests. Sample is the same as in Fig. 2, error bars show 99% confidence limits, and the dashed line shows the $\alpha = 0.05$ level. "Jitter" includes 1000 repetitions (rectangular window, $W = 11$ ms) and the partially hollowed windows are with HF of 0.42 (rectangular); 0.6 (Gaussian); and 0.63 (triangular). In contrast to the other methods, the modified convolution method based on a partially hollowed window is unbiased. (B) Power of several tests. Sample is the same as in Fig. 5B; other conventions are the same as in (A). The partially hollowed windows have higher power – and equivalently a lower miss rate – than the full windows.

3.4. Application of the modified convolution method to real spike trains

Correct empirical assessment of any statistical method relies on the data employed for testing it. For testing false positive rates, the data must conform exactly to the null hypothesis. Up to now the assessment of the jittering, convolution, and modified convolution methods was based on pairs of artificial spike trains. These trains were generated independently, with or without slow temporal correlations (Eq. (2)), and then synchronous spikes were inserted into both trains (Section 2.2.1; Grün et al., 1999). The latter approach was employed because then trains could be generated using predetermined firing rates, various sample sizes, and known levels of correlation and synchrony.

However, the generation of artificial spike trains necessarily requires assuming some underlying probability distribution such as Poisson or gamma. The above results (and specifically the values of the optimal HFs for the various window types) were consistent for Poisson and 2nd, 3rd, and 4th order gamma spike trains (data not shown). However, cortical neurons may fire according to time-varying, mixed, or other distribution functions. Thus, we next used real spike trains of well-isolated single-units recorded from the premotor cortices of monkeys that participated in a task requiring memorization of the forthcoming reach direction and grasp type followed by actual prehension.

To assess the false positive rate of the various methods using pairs of simultaneously recorded spike trains, we shuffled the trial order of one spike train relative to the other, for each reach-grasp combination separately (conditionally independent (CI) shuffling; Stark et al., 2008). This results in un-synchronized pairs of spike trains but maintains the original realistic statistical features of each train. We applied the jittering, convolution, and modified convolution methods to these data, which consisted of 629 spike train pairs, each recoded over a total of 459 trials (median; range: 375–581); the median number of spikes per train was 2648 (range: 518–26,097). As for the artificial data (Fig. 6), the jittering and convolution methods yielded biased results whereas the modified convolution method yielded unbiased results (Fig. 7A). This indicates that our results are independent of a specific statistical structure (e.g. a time-varying Poisson or gamma process) of the spike trains hence generalizing the optimality of the modified convolution method.

We then applied the modified convolution method to the original (non-shuffled) spike train pairs. In this case, the test power cannot be determined because the fraction of synchronized pairs is unknown. In fact, the motivation for developing an unbiased and sensitive test was exactly to enable accurate detection of pairwise synchrony in real spike trains. For the premotor cortex data, we found that 49/629 (7.8%) of the spike train pairs were syn-

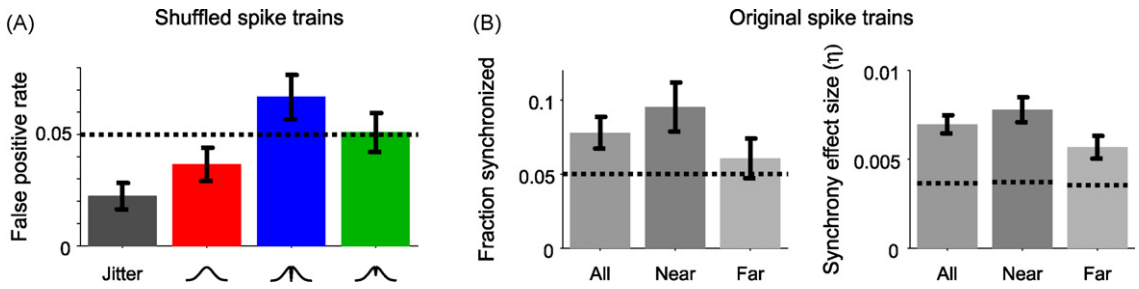

Fig. 7. Analysis of real spike trains. (A) False positive rate for spike trains of well-isolated SUs recorded from the premotor cortices of monkeys during a prehension task. Sample contains 629 SU pairs. Spike trains were shuffled one relative to the other so pairs are unsynchronized (see Section 3.4); under these circumstances, the false positive rate of an unbiased test should be exactly α (here, 0.05). Error bars show binomial S.E. The jittering and convolution methods are conservative; the use of a hollowed window is permissive; and the modified convolution method employing a partially hollowed Gaussian window (HF = 0.6) is unbiased. (B) Synchrony among SU pairs. Sample is the same as in (A), except that the spike trains were not shuffled so that any synchrony in the original data is retained. *Left*: fraction of synchronized SU pairs, partitioned according to inter-unit distance. Sample sizes are 629 (all), 315 (near), and 314 (far). *Right*: synchrony effect sizes (η) for the subset of synchronized SU pairs (all/near/far: 49/30/19 pairs). Error bars show S.E. and dashed lines show the effect sizes expected by chance, estimated for each synchronized pair using Eq. (6) and averaged over all pairs.

chronized, a fraction significantly higher than the chance level of $\alpha = 0.05$ (binomial test: $p = 0.0017$; Fig. 3B, *left*). Among the synchronized pairs, the median η was 0.0062, significantly higher than the corresponding median η_{chance} among the same pairs (0.003; Mann–Whitney *U*-test: $p \ll 0.001$; Fig. 7B, *right*). Moreover, both the probability of spike trains to be synchronized and the synchrony effect size depended on inter-electrode distance, being as high as 9.5% and 0.0065, respectively, for spike train pairs recorded from different electrodes in the same premotor region (PMd or PMv; median inter-electrode distance, 0.92 mm), and as low as 6.1% and 0.0049 for different-region pairs (median inter-electrode distance, 5.3 mm; Fig. 7B).

4. Discussion

4.1. Properties of the jittering and straight-forward convolution methods

For testing the null hypothesis of "no precise temporal correlations between two spike trains", the jittering method and a straight-forward application of the convolution method yield equivalent results (Fig. 1). In particular, both methods allow testing the same null hypothesis and enable flexibility in the definition of temporal precision. Thus, neither method is confounded by slow temporal correlations, trial-to-trial variability, or other sources of inter- or intra-trial non-stationarity (Brody, 1999; Ben-Shaul et al., 2001).

However, the analogy between the two methods explains why, for uncorrelated data, the *p*-value distribution yielded by the jittering method is highly non-uniform: the distribution is essentially identical to the distribution obtained from the application of the straight-forward convolution method (Fig. 2). The latter distribution is non-uniform because of two reasons. First, CCH counts are discrete, whereas proper probability estimation requires an integral transformation of a continuous random variable (Pearson, 1950). Second, the estimation of the expected values in a CCH is biased due to over-fitting the raw data. Thus, we suggest caution in the use of the straight-forward convolution and jittering methods.

4.2. Properties of the modified convolution method

The jittering method and the straight-forward convolution method are both conservative and have relatively low power. In contrast, the modified convolution method, using a partially hollowed window with a continuity correction, yields unbiased results (Fig. 4) and has higher power (Fig. 5). Statistical power depends on several factors including effect and sample sizes; single cortical neurons, recorded over several minutes, typically provide a few thousand spikes. In such data, the modified convolution method enables detecting weak synchrony with a negligible miss rate (less than 1%; Fig. 5C).

Beyond the superior statistical properties of the modified convolution method over the jittering method, there are practical considerations that favor the use of the former. Following the computation of a raw CCH, the jittering method requires hundreds to thousands of resampling repetitions before estimating probabilities and thus relatively long computing time. In contrast, the convolution method does not require resampling the spike trains or repetitive CCH computations and is therefore computationally faster by several (2 or more) orders of magnitude.

There are a couple of limitations of the modified convolution method that deserve attention. First, when firing rates are high, ISIs are small and convolution is not identical to jittering (Fig. 1G). This is because jittering is carried out with special care not to violate certain statistical features of each spike train (Section 2.1.3). In such high firing rate cases, each spike train should be diluted prior to computation (Section 2.1.1). Quantitatively, if the minimal ISI to be left after jittering is R and the window width is W, each spike train should be diluted until the fraction of ISIs smaller than $W+R$ is negligible.

Second, convolution (or, for that matter, jittering) with any window, partially hollowed or full, always flattens peaks. Therefore, slow temporal CCH features may in principle turn out significant following convolution even when a narrow window is employed. To illustrate this problem, consider a triangular CCH peaking at zero-lag (1 ms bins) and reducing gradually to zero at time lags of ± 100 ms. In that case, a count of more than 1000 coincidences is required to cross a significance threshold of 0.05. Using stationary spike trains with 5 spikes/s and rearranging Eq. (5) to yield $T = \lambda/(F^2 B)$ shows that 1000 coincidences would be accumulated after ~40,000 s of recording. Thus, although false detections of precise temporal correlations due slow temporal correlations are a concern, they are still largely a theoretical case since such enormous amounts of data are not readily available from extra-cellular recordings. If however the suspicion does arise, then the sample size (CCH count) required to cross a given significance threshold can be computed and this potential confound excluded.

4.3. Extensions and applications

The modified convolution method using a partially hollowed window is adequate for estimating precise temporal correlations in spike train pairs and thus jittering is unnecessary for that purpose. However, it may be desired to jitter spikes for other purposes requiring multiple realizations of a given spike train under the null hypothesis of no precise timing. As illustrated here, straight-

forward jittering within a full (rectangular or other) window is inadequate. For jittering two spike trains an unbiased jittering window does not exist since a narrowly and partially hollowed window cannot be de-convolved into two identical windows. However, in some applications of the jittering method, only a single spike train was manipulated (Hatsopoulos et al., 2003; Pazienti et al., 2007). Specifically when precise temporal correlations between two spike trains are to be assessed, unbiased jittering can be carried out by jittering only one train using a partially hollowed rectangular window, yielding results equivalent to convolving the CCH with the same window. In general, since any jittering within a full window over-fits and jittering within a hollowed window under-fits the raw spike trains, a partially hollowed window should be used in any jittering application.

The unitary-event analysis method has been developed (Grün et al., 1999) and applied (Riehle et al., 1997; Maldonado et al., 2008) for detecting synchronous and near-synchronous activity of multiple spiking neurons. The method enables time-resolved detection of above-chance synchronous events, thereby yielding, for two spike trains, more detailed information than the above-chance coincident count in the zero-lag bin of a CCH (or the diagonal of a joint peri-stimulus time histogram; Aertsen et al., 1989). The estimation of chance level synchrony in the straight-forward method (Grün et al., 2002a) rests on the unrealistic assumptions of within-trial and inter-trial stationarity. Thus, to account for within-trial non-stationarity, the statistical evaluation of chance-level has been modified to employ small time windows (Grün et al., 2002b). To account for trial-to-trial variability as well, convolution methods were employed (Pauluis and Baker, 2000; Maldonado et al., 2008). As shown here, convolution with a full window is over-conservative and therefore convolution with a partially hollowed window is recommended in the specific case of unitary-event analysis of two spike trains.

4.4. Future directions

While the current work focused on interactions between two spike trains, the modified convolution method can be easily extended to the study of three- (or more) neuron interactions. For three (or N) spike trains, the one-dimensional convolution window used for pairs of spike trains should be replaced by a 2D (or $N-1$ dimensional) kernel with a partially hollowed central bin. From that point on computations should proceed exactly as for the pair-wise case. Yet the exact HF for various types of N-dimensional kernels remains to be determined using, for instance, Monte-Carlo simulations as in the present work.

The CCH time lag of interest depends on the coding model one has in mind. Here, we focused on coincident spiking (precise temporal synchrony), consistent with common input or recurrent circuits, and ignored the non-zero time lags expected for direct (mono- or poly-synaptic) circuits. The mathematical analysis of non-zero lag bins is identical to the analysis of synchrony with one technical exception: when sample size is finite, there are potentially less data for the non-zero lag bins. By trimming spike train tails and flipping one of the CCHs (Section 2.1.1), this is no longer a consideration and thus the modified convolution method developed here can be applied to any single time lag.

However, when the Synfire chain model (Abeles, 1982a) is considered, one expects precise temporal correlations to appear at multiple time lags other than zero. While recent work has directly addressed the issue of detecting traces of Synfire chain activity in large neural networks (Schrader et al., 2008), there are other connectivity patterns that may give rise to non-zero time lags. These considerations raise the issue of multiple comparisons in a CCH, not addressed in the current work. In such a case, the significance threshold (α level) should be reduced according to the number of independent comparisons. One way is to use a standard correction for multiple comparisons such as the Tukey-Cramer or the Bonferroni correction (Hochberg and Tamhane, 1987). However, these corrections are typically conservative, depend on the number of time lags actually used for CCH computation, and ignore possible correlations between adjacent CCH bins. Thus, while the modified convolution method can be directly applied to the study of precise temporal synchrony, future work should rigorously address the issue of multiple comparisons in a CCH.

Acknowledgements

We thank Alit Stark for critical comments. This research was supported in part by a Center of Excellence grant (1564/04) administered by Israel Science Foundation and the Deutsch-Israelische Projectkooperation.

References

Abeles M. Local cortical circuits: an electrophysiological study. Berlin: Springer–Verlag; 1982a.
Abeles M. Quantification, smoothing, and confidence limits for single-units' histograms. J Neurosci Methods 1982b;5:317–25.
Abeles M. The quantification and graphic display of correlations among three spike trains. IEEE Trans Biomed Eng 1983;30:235–9.
Abeles M, Gat I. Detecting precise firing sequences in experimental data. J Neurosci Methods 2001;107, 141–54;112: 203.
Aertsen AM, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of "effective connectivity". J Neurophysiol 1989;61:900–17.
Baker SN, Gerstein GL. Improvements to the sensitivity of gravitational clustering for multiple neuron recordings. Neural Comput 2000;12:2597–620.
Baker SN, Spinks R, Jackson A, Lemon RN. Synchronization in monkey motor cortex during a precision grip task. I. Task-dependent modulation in single-unit synchrony. J Neurophysiol 2001;85:869–85.
Ben-Shaul Y, Bergman H, Ritov Y, Abeles M. Trial to trial variability in either stimulus or action causes apparent correlation and synchrony in neuronal activity. J Neurosci Methods 2001;111:99–110.
Brody CD. Correlations without synchrony. Neural Comput 1999;11:1537–51.
Date A, Bienenstock E, Geman S. On the temporal resolution of neural activity. Technical report, Division of Applied Mathematics, Brown University (http://www.dam.brown.edu/people/elie/papers/temp–res.ps), 1998.
Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. Nat Neurosci 2008;11:823–33.
Gerstein GL. Searching for significance in spatio-temporal firing patterns. Acta Neurobiol Exp (Wars) 2004;64:203–7.
Grün S, Diesmann M, Grammont F, Riehle A, Aertsen A. Detecting unitary events without discretization of time. J Neurosci Methods 1999;94:67–79.
Grün S, Diesmann M, Aertsen A. Unitary events in multiple single-neuron spiking activity. I. Detection and significance. Neural Comput 2002a;14:43–80.
Grün S, Diesmann M, Aertsen A. Unitary events in multiple single-neuron spiking activity. II. Nonstationary data. Neural Comput 2002b;14:81–119.
Hatsopoulos N, Geman S, Amarasingham A, Bienenstock E. At what time scale does the nervous system operate? Neurocomputing 2003;52–54:25–9.
Hochberg Y, Tamhane AC. Multiple comparison procedures. New York: Wiley; 1987.
Maldonado P, Babul C, Singer W, Rodriguez E, Berger D, Grün S. Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. J Neurophysiol 2008;100:1523–32.
Neyman J, Pearson ES. The testing of statistical hypotheses in relation to probabilities a priori. Proc Camb Phil Soc 1933; 29:492–510; Reprinted in: Neyman J, Pearson ES. Joint statistical papers. Cambridge: Cambridge University Press; 1967, pp. 186–202.
Palm G, Aertsen AM, Gerstein GL. On the significance of correlations among neuronal spike trains. Biol Cybern 1988;59:1–11.
Pazienti A, Diesmann M, Grün S. Bounds of the ability to destroy precise coincidences by spike dithering. In: Advances in brain, vision, and artificial intelligence, lecture notes in computer science. Berlin Heidelberg: Springer–Verlag; 2007. pp. 428–37.
Pauluis Q, Baker SN. An accurate measure of the instantaneous discharge probability, with application to unitary joint-even analysis. Neural Comput 2000;12: 647–69.
Pearson ES. On questions raised by the combination of tests based on discontinuous distributions. Biometrika 1950;37:383–98.
Perkel DH, Gerstein GL, Moore GP. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. Biophys J 1967;7:419–40.
Perkel DH, Gerstein GL, Smith MS, Tatton WG. Nerve-impulse patterns: a quantitative display technique for three neurons. Brain Res 1975;100:271–96.
Prut Y, Vaadia E, Bergman H, Haalman I, Slovin H, Abeles M. Spatiotemporal structure of cortical activity: properties and behavioral relevance. J Neurophysiol 1998;79:2857–74.

Riehle A, Grün S, Diesmann M, Aertsen A. Spike synchronization and rate modulation differentially involved in motor cortical function. Science 1997;278:1950–3.
Schrader S, Grün S, Diesmann M, Gerstein GL. Detecting synfire chain activity using massively parallel spike train recording. J Neurophysiol 2008;100:2165–76.
Shmiel T, Drori R, Shmiel O, Ben-Shaul Y, Nadasdy Z, Shemesh M, Teicher M, Abeles M. Neurons of the cerebral cortex exhibit precise interspike timing in correspondence to behaviour. Proc Natl Acad Sci USA 2005;102:18655–7.
Stark E, Asher I, Abeles M. Encoding of reach and grasp by single neurons in premotor cortex is independent of recording site. J Neurophysiol 2007;97:3351–64.
Stark E, Globerson A, Asher I, Abeles M. Correlations between groups of premotor neurons carry information about prehension. J Neurosci 2008;28:10618–30.