

Correlations without synchrony

Carlos D. Brody¹

Computation and Neural Systems Program, Caltech, Pasadena CA 91125, USA. email: carlos@cns.caltech.edu

Accepted for publication in *Neural Computation*, 25 November 1998; published as:

Brody, C.D., *Neural Computation*, **11**:1537–1551 (1999).

Abstract

Peaks in spike train correlograms are usually taken as indicative of spike timing synchronization between neurons. Strictly speaking, however, a peak merely indicates that the two spike trains were not independent. Two biologically-plausible ways of departing from independence which are capable of generating peaks very similar to spike timing peaks are described here: covariations over trials in response *latency* and covariations over trials in neuronal *excitability*. Since peaks due to these interactions can be similar to spike timing peaks, interpreting a correlogram may be a problem with ambiguous solutions. What peak shapes do latency or excitability interactions generate? When are they similar to spike timing peaks? When can they be ruled out from having caused an observed correlogram peak? These are the questions addressed here. A companion paper (Brody, 1998) proposes quantitative methods to tell cases apart when latency or excitability covariations cannot be ruled out.

1 Introduction

Suppose that the spike trains of two neurons, recorded simultaneously during many identically prepared experimental trials, have been obtained. A standard method to assess the presence of interactions between the spike trains – beyond those expected by chance given each neuron’s peri-stimulus time histogram (PSTH) – is to compute their shuffle-corrected cross-correlogram (Perkel et al., 1967; Palm et al., 1988; Aertsen et al., 1989). The name *shuffle-corrected cross-correlogram* will be henceforth abbreviated to *crosscovariogram* or simply *covariogram*.² Peaks in covariograms are usually interpreted as signaling the presence of spike timing synchronization between the two neurons. Strictly speaking, however, a peak in a covariogram merely indicates that the two spike trains were not independent, and synchronizing the spike times of the two neurons is only one of many possible ways to depart from independence. Figure 1 shows three very different ways to depart from

¹Present address: Instituto de Fisiología Celular, UNAM, Apdo. Postal 70-253, México D.F. 04510, México.

²The abbreviation *covariogram* comes from the fact that the computation of the shuffle-corrected cross-correlogram is exactly analogous to the computation of covariance when the variables of interest are scalars rather than spike trains (Aertsen et al., 1989; Brody, 1997a).

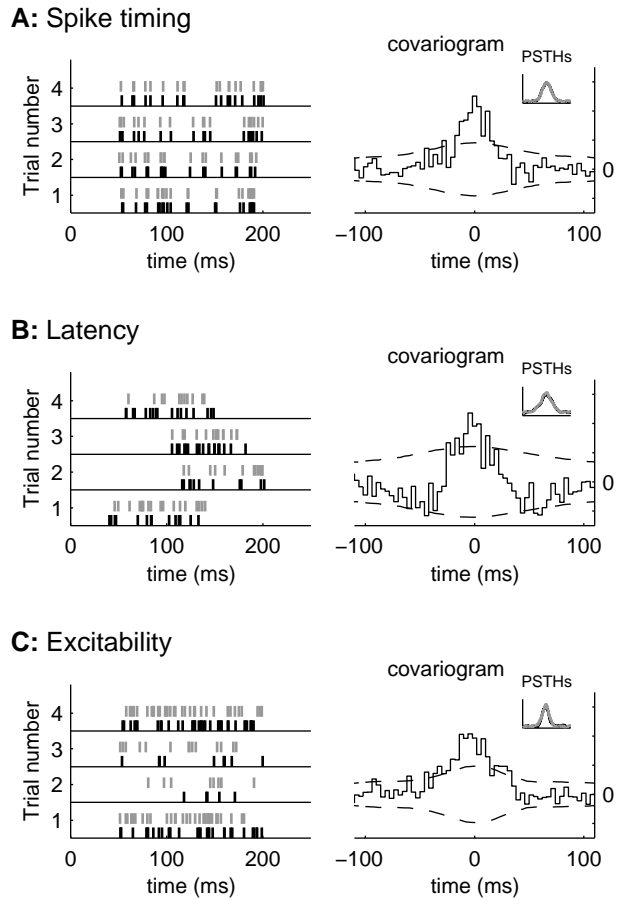


Figure 1: Three types of covariations. Despite being very different, all three shuffle-corrected correlograms (henceforth called *covariograms*) look very similar. Each row illustrates a type of covariation: On the left is a raster plot of two simulated cells, and on the right is the covariogram of spike trains made in a similar fashion. (Parameters used for rasters on the left were set to extreme values, to emphasize illustrative clarity; parameters used for covariograms on the right were set to “physiologically plausible” values.) **A:** On each trial, most spikes in cell 1 have a corresponding, closely timed, spike in cell 2. Both cells have the same response latency and overall firing rate in all trials. This will be called a *spike timing* covariation. **B:** Spikes in cell 1 do not have a corresponding spike in cell 2: on each trial, the two spike trains were generated independently of each other. But the overall *latency* of the response varies together over trials. (The word *latency* will be used here to indicate the time shift of the whole response, not just of the first spike.) **C:** Again, on each trial the spikes for the two cells were generated independently of each other. But the total magnitude of the response (which will be called the *excitability*) varies together over trials. Zero counts on the covariogram y axes is the expected value if the two cells are independent; the dashed lines are significance limits. The inset at the top right of each covariogram shows the PSTHs of the two cells involved, plotted on axes that are 250 ms wide and 60 Hz tall.

independence, all of which lead to very similar covariograms. Despite their similarity, each case should be interpreted very differently, both in terms of the mechanisms that could cause it and in terms of its functional significance. All three types of covaria-

tions illustrated (which will be called *spike timing*, or *latency*, or *excitability* covariations) are biologically plausible. Thus, being aware of the different possibilities, and disambiguating them, is important.

This paper will explain how latency and excitability covariations lead to a peak in the covariogram (spike timing covariations have been treated before, e.g. Perkel et al. 1967). In doing so, the paper will also explain under what conditions their peaks are similar to peaks caused by spike timing covariations. Rules of thumb for being alert to the possibility of ambiguous covariograms will be emphasized; a companion paper (see Brody, 1998) describes more quantitative methods, which attempt to dispel the ambiguity when it arises. A preliminary version of the results presented here has appeared in abstract form (Brody, 1997b).

2 Notation and correlogram methods

The spike trains of two cells will be represented by two time-dependent functions, $S_1(t)$ and $S_2(t)$. They will be assumed *binned* and collected over many identically prepared experimental trials, indexed by a superscript r . For times outside the r -th trial, $S_1^r(t)$ will be *defined* to be zero, and similarly for $S_2^r(t)$. The cross-correlogram of each trial is then

$$C^r(\tau) \equiv \sum_{t=-\infty}^{\infty} S_1^r(t) S_2^r(t + \tau) \equiv S_1^r \odot S_2^r. \quad (1)$$

Let $\langle \rangle$ represent averaging over trials r , and define $P_i(t) \equiv \langle S_i^r(t) \rangle$. If spike times are measured relative to a stimulus, this is the PSTH of S_i . The covariogram of S_1 and S_2 is then defined as

$$\begin{aligned} V &\equiv \langle (S_1^r - P_1) \odot (S_2^r - P_2) \rangle \\ &= \langle S_1^r \odot S_2^r \rangle - P_1 \odot P_2. \end{aligned} \quad (2)$$

The two terms in equation (2) are known as the *raw cross-correlogram* $R = \langle S_1^r \odot S_2^r \rangle$ and the *shuffle corrector*³ $K = P_1 \odot P_2$. If S_1 and S_2 are independent, then the expected value of V is zero:

$$\begin{aligned} E\{V\} &= E\{(S_1^r - P_1) \odot (S_2^r - P_2)\} \\ &= E\{S_1^r - P_1\} \odot E\{S_2^r - P_2\} = 0. \end{aligned} \quad (3)$$

Therefore, significant departures of V from zero indicate that the two cells were not independent, regardless of what distributions that S_1^r and S_2^r were drawn from. Estimating the significance of departures of V from 0 requires some assumptions: for

³The *shift predictor* (Perkel et al., 1967) is very similar to the shuffle corrector, except that K is replaced by $D = \langle S_1^r \odot S_2^{\Pi(r)} \rangle$, where $\Pi(r)$ is some permutation of the stimulus presentations r (and the corresponding substitution is made in equation (2)). If different trials are independent of one another, then the expected value of the shift predictor D is equal to the expected value of the shuffle corrector K . Thus they are both estimators of the same function. In practice it is preferable to use K instead of D since the former is a less noisy estimator: K can be written as the average of D , taken over the set of all possible permutations Π (Palm et al., 1988).

the null hypothesis, it will be assumed that S_1 is independent of S_2 , that different trials of S_1 are independent of each other, and that different bins within each trial of S_1 are independent of each other (similar assumptions for the trials and bins of S_2 will also be made). If $P_i(t)$ and $\sigma_i^2(t)$ are the mean and variance of $S_i^r(t)$ over trials r , and N_{trials} is the number of trials in the experiment, then the variance in the null hypothesis for V is

$$\sigma_V^2(t) = (\sigma_1^2 \odot \sigma_2^2 + P_1^2 \odot \sigma_2^2 + \sigma_1^2 \odot P_2^2) / N_{\text{trials}}. \quad (4)$$

In practice, one uses the sample means and variances to calculate $\sigma(t)$; the 2σ limits, calculated in this way, are displayed as dashed lines in the covariograms throughout this paper.⁴ While σ is a general measure of the spread of a distribution, more assumptions must be made in order to use it to assign a specific number to a significance limit: e.g., if the distribution is assumed Gaussian then 2σ represents the 95% confidence limit. No particular assumption will be made here.

Joint Peristimulus Time Histograms, or JPSTHs (Aertsen et al., 1989), will also be used. The unnormalized JPSTH is a matrix of covariances with elements defined as

$$J(t_1, t_2) = \langle S_1^r(t_1) S_2^r(t_2) \rangle - \langle S_1^r(t_1) \rangle \langle S_2^r(t_2) \rangle, \quad (5)$$

while the normalized JPSTH is a matrix of correlation coefficients with elements defined as

$$J_N(t_1, t_2) = \frac{J(t_1, t_2)}{\sigma_1(t_1) \sigma_2(t_2)}, \quad (6)$$

If $S_1(t_1)$ and $S_2(t_2)$ are independent then the expected values of $J(t_1, t_2)$ and $J_N(t_1, t_2)$ are zero. Correlation coefficients are bounded within $[-1, 1]$. If $J_N(t_1, t_2) = 1$, then $S_1(t_1)$ and $S_2(t_2)$ are perfectly correlated (that is, $S_1(t_1) = \alpha S_2(t_2)$ for some positive constant α), while if $J_N(t_1, t_2) = -1$, then $S_1(t_1)$ and $S_2(t_2)$ are perfectly *anticorrelated* (that is, $S_1(t_1) = -\alpha S_2(t_2)$). The JPSTHs displayed in the figures of this paper are all normalized JPSTHs.

The covariogram V can be obtained from the unnormalized JPSTH J by summing along t_1 while keeping $\tau = t_2 - t_1$ constant.

3 What covariogram shapes do latency and excitability covariations generate?

3.1 Latency covariations

Take the responses of two independent neurons. Since they are independent, their covariogram is zero (within sampling noise);

⁴To see the where equation (4) comes from, consider two *independent* scalars x and y with means p_x and p_y and variances σ_x^2 and σ_y^2 , respectively. The variance of their product is $E\{x^2\}E\{y^2\} - E\{x\}^2 E\{y\}^2 = (\sigma_x^2 + p_x^2)(\sigma_y^2 + p_y^2) - p_x^2 p_y^2 = \sigma_x^2 \sigma_y^2 + p_x^2 \sigma_y^2 + \sigma_x^2 p_y^2$. Equation (4) is analogous. The factor of N_{trials} comes from averaging over trials.

hence the raw cross-correlogram and the shuffle corrector are approximately equal:

$$\underbrace{V}_{\text{covariogram}} = \underbrace{R}_{\text{raw x-corr}} - \underbrace{K}_{\text{corrector}} \approx 0 \implies K \approx R. \quad (7)$$

Now for each trial r , take the responses of both neurons and shift both their spike trains, together, by some amount of time t^r (the shift time t^r should be different for different trials). This type of interaction between the neurons is dubbed here a *latency covariation*. How will it affect V ?

Let us ask how it affects each of the two terms of V , that is, R and K . The raw correlogram R will not be affected, since it only depends on relative spike times between the two neurons (see equation (1)), and on each trial both spike trains were shifted together. In contrast, the shuffle corrector K will be affected. It is the correlogram of the two PSTHs, and the PSTHs are broadened by the temporal jitter introduced by the shifts t^r . Thus K is broader than before the latency shifts. Since the total number of spikes remains the same, the integral of the PSTHs will not have changed; the integral of K , too, will not have changed. In summary, the latency shifts will make K broader, and therefore shallower, while having no effect on R . Panel A of Figure 2 shows a schematic of how subtracting the broadened, shallower, K from R leaves a peak in R outstanding in V . The peak is flanked by slight negative dips. The most important point to notice about this schematic is that the width and shape of the peak in V is largely determined by the width and shape of the peak in R .

Unless the latency shifts are very large, the width of the peak in R , and hence in V , will be smaller but of the same order of magnitude as the width of the peak in K , which in turn is determined by the width of peaks in the cell’s PSTHs. Panels B–F in Figure 2 show a numerical experiment illustrating latency covariations. The covariogram peak width is ≈ 50 ms, while PSTH peak widths are ≈ 100 ms. For the simple Poisson-like processes used here, and for symmetrical cells, the *autocovariograms* of each cell (panels E and F) have a shape similar to the *crosscovariogram* of the two cells (panel C).

3.2 Excitability covariations

Consider a cell whose response can be characterized as the sum of a stimulus-induced response plus a background firing rate. Let us write this in terms of firing rates⁵ as

$$\underbrace{F^r(t)}_{\text{Firing rate during trial } r} = \underbrace{\zeta^r Z(t)}_{\text{Stimulus-induced}} + \underbrace{\beta^r B}_{\text{Background}} \quad (8)$$

Here $Z(t)$ is the typical stimulus-induced firing rate, B is a constant function over the time of a trial, representing the typical

⁵The description given in equation (8) amounts to describing the cell with a generative model; but all that is being specified about the model is the expected value of its response on each trial. Thus, if $M^r(t)$ is the model’s response during trial r , then $F^r(t) = E\{M^r(t)\}$. Note that the expectation here is *not* taken *across* trials, but is the expected response for a *single* trial. Think of this as fixing the model’s parameters at values appropriate for trial r , and averaging over many runs of the model at those parameters.

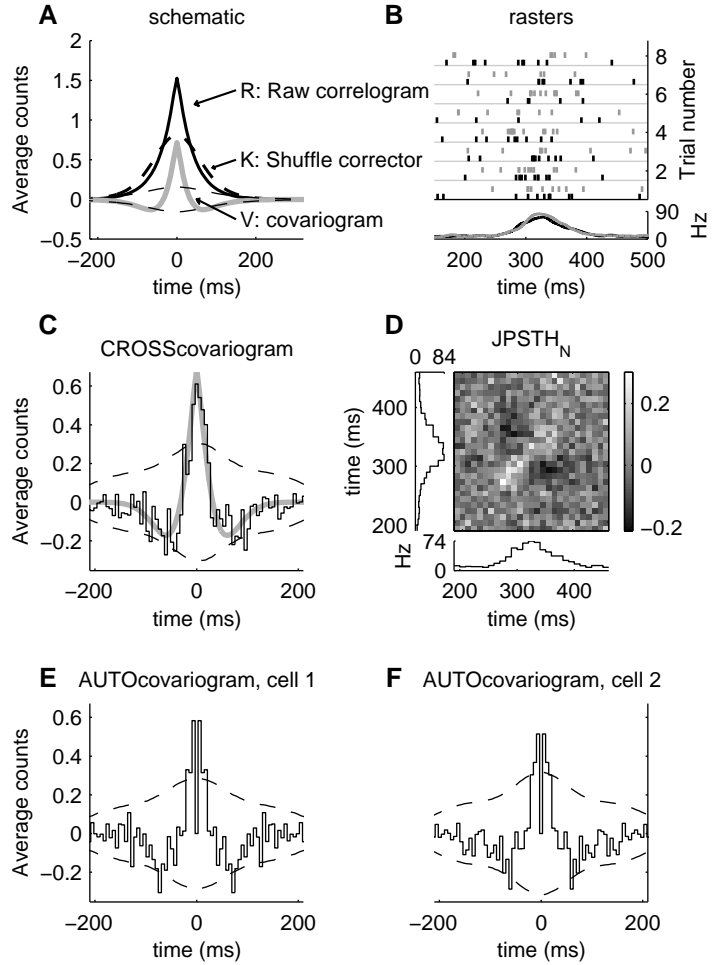


Figure 2: Latency covariations. **A:** Schematic of how latency covariations lead to a peaked covariogram (see text for explanation). **B:** Eight out of two hundred artificial rasters used to illustrate latency covariations. Below the rasters are the smoothed PSTHs of both cells. The spike trains were made by simulating two *independent* Poisson cells, each raster pair of which was then shifted together in time by a random amount drawn anew for each trial from a Gaussian distribution with mean 0 ms and standard deviation 15 ms. [Details: The time-varying firing rate from which Poisson events for each cell were generated, before the time shifts, was $(100\text{Hz}) \exp((t - 100)^2 / (2 \cdot 40^2))$ for $t > 100$, zero otherwise, with t in milliseconds. After the time shifts, independent events at a rate of 10 Hz were added to each cell to represent background firing.] **C:** Covariogram of the two cells. The thick grey line is the analytical expected value of the covariogram, computed with knowledge of the parameters and procedures used generate the spike trains. Dashed lines are significance limits. **D:** Normalised JPSTH of the two cells, flanked by their PSTHs. Notice the diagonal peak and weak but present off-diagonal troughs. **E, F:** Autocovariograms of each of the two cells. These are computed following equation (2), except one cell is used, not two, and the central bin ($\tau = 0$), which for autocovariograms is much larger than any other, has been arbitrarily set to zero here for display purposes.

background firing rate, and two “gain” factors, ζ^r and β^r , have been included to represent possible changes in the state of the cell (e.g. changes over trials in the resting potential of the cell

(Carandini and Ferster, 1997)). The gain factors ζ^r and β^r will be allowed to be different for different trials. Two assumptions are being made here: (1) state changes are slower than the time of a single trial; (2) the greatest effect of state changes is on the magnitude of the background and stimulus-induced rates, not on their temporal shape. These assumptions allow factoring out the effect of state changes into the scalar gain factors ζ^r and β^r .

ogram V ? In what follows, write the covariance of two scalars a and b as $\text{cov}(a, b) \equiv \langle ab \rangle - \langle a \rangle \langle b \rangle$, and drop the superscripts r for legibility. Using equations (2) and (8), and the fact that the gain parameters factor out,

$$V = \overbrace{\text{cov}(\zeta_1, \zeta_2)}^{\text{amplitude}} \overbrace{Z_1 \odot Z_2}^{\text{shape}} + \text{cov}(\zeta_1, \beta_2) Z_1 \odot B_2 + \text{cov}(\beta_1, \zeta_2) B_1 \odot Z_2 + \text{cov}(\beta_1, \beta_2) B_1 \odot B_2. \quad (9)$$

Similarly, the JPSTH (before normalization) is

$$J(t_1, t_2) = \text{cov}(\zeta_1, \zeta_2) Z_1(t_1)Z_2(t_2) + \text{cov}(\zeta_1, \beta_2) Z_1(t_1)B_2 + \text{cov}(\beta_1, \zeta_2) B_1Z_2(t_2) + \text{cov}(\beta_1, \beta_2) B_1B_2, \quad (10)$$

where the time-dependence of B_1 and B_2 has been dropped from the notation since they are constant functions.

When the stimulus-induced firing rate is much greater than the background firing rate, the first term in equation (9) is the dominant one. The *shape* of V will then be given by $Z_1 \odot Z_2$ (in this limiting case, this is also the shape of the corrector K , which has a width determined by the width of peaks in the cell's PSTHs), while the *amplitude* of V will be given by $\text{cov}(\zeta_1, \zeta_2)$. A similar point has been made by Friston (1995), whose work will be discussed further in a companion paper (Brody 1998; see also Vaadia et al. 1995). Figure 3 shows a numerical experiment illustrating excitability covariations. For the simple Poisson-like processes used here, and for symmetrical cells, the *autocovariograms* of the cells (panel D) are similar to the *crosscovariogram* (panel B), much as was the case with latency covariations (Figure 2).

An easily computable and tell-tale measure of excitability covariations is the integral (i.e. sum) of the covariogram, since it is proportional to the covariation in the mean firing rates of the two cells: $\sum_{\tau} V(\tau) = \text{cov}(n_1^r, n_2^r)$, where n_i^r is the total number of spikes fired by cell i during trial r . For completeness, the proof follows:

$$\begin{aligned} \sum_{\tau=-\infty}^{\infty} C^r(\tau) &= \sum_{\tau=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} S_1^r(t)S_2^r(t+\tau) \\ &= \sum_{p,q} S_1^r(p)S_2^r(q) \\ &= \sum_p S_1^r(p) \sum_q S_2^r(q) \\ &= n_1^r n_2^r. \end{aligned} \quad (11)$$

Thus $\sum_{\tau} R(\tau) = \langle n_1^r n_2^r \rangle$. Similarly, $\sum_{\tau} K(\tau) = \langle n_1^r \rangle \langle n_2^r \rangle$. Hence

$$\sum_{\tau} V(\tau) = \langle n_1^r n_2^r \rangle - \langle n_1^r \rangle \langle n_2^r \rangle = \text{cov}(n_1, n_2) \quad (12)$$

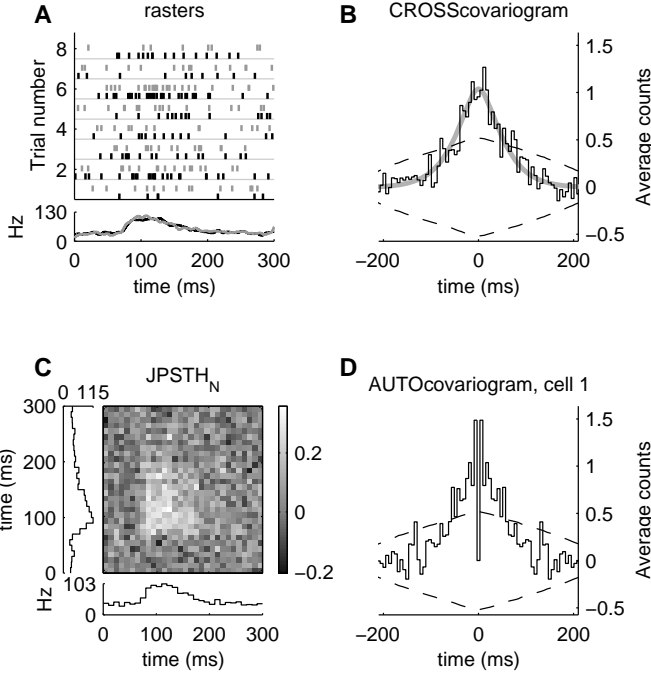


Figure 3: Excitability covariations. **A:** Eight out of two hundred rasters, made by simulating two independent Poisson cells with covarying gains ζ_1 and ζ_2 (see text). Both gains were set to be equal to each other on each trial, and were a random number drawn anew for each trial from a Gaussian with mean 1 and standard deviation 1 (negative gains were set to zero). Below the rasters are the smoothed PSTHs of the two cells. [Details: $Z_1(t)$ and $Z_2(t)$ were both set to be, before multiplying by the gains, $(70 \text{ Hz}) \cdot ((t - 70)/30) \cdot \exp((100 - t)/30)$ if $t > 70$, zero otherwise, with t in milliseconds. After multiplying by the gain, a constant rate (same for all trials) of 35 Hz was added to represent background firing.] **B:** Covariogram of the two cells. The thick grey line is the analytical expected value, computed from equation (9) with knowledge of the parameters used to generate the spike trains. Dashed lines are significance limits. Notice that the width of the peak is comparable to twice the width of the peak in the PSTHs. In the example shown here, the background firing rate is not negligible, so the covariogram doesn't quite follow the shape of the shuffle corrector (which is not shown), but follows the shape of $Z_1 \odot Z_2$, the "stimulus-induced" parts of the PSTHs. **C:** Normalised JPSTH of the two cells, flanked by their PSTHs. **D:** Autocovariogram of one of the cells; the other is similar. The central bin ($\tau = 0$) has been set to zero for display purposes.

Now take two cells, indexed by the subscripts 1 and 2, with responses characterized as in equation (8). Suppose their only interaction is through their gain parameters. This has been dubbed here an *excitability* covariation. What is their covari-

3.3 Spike timing covariations

Figure 4 shows a numerical experiment illustrating spike timing covariations. Three major points in comparison to latency and excitability covariations are that: (1) for the simple Poisson-like processes used here, where there was no burstiness and the spike timing interaction was between individual spikes of the two cells, the *autocovariograms* are flat and not at all similar to the *crosscovariogram* of the two cells. This is in contrast to the latency or excitability covariations cases, and allows using the autocovariograms as a first test to distinguish spike timing from latency or excitability covariations. (2) While latency and excitability covariations involve coordination between as little as a single parameter of the two cells on each trial (overall latency in one case, gain in the other), spike timing covariations will typically involve coordination between many parameters on each trial (many individual spike times). (3) Given arbitrary network connectivities, spike timing covariogram shapes are much more arbitrary than latency or excitability covariogram shapes. In particular, while the latter are tied to the shapes of the PSTHs, the former are not.

4 Discussion

Peaks in spike train covariograms are typically interpreted as evidence of spike timing synchronization, but other ways to depart from independence can generate covariogram peaks very similar to spike synchronization peaks. Two such departures have been described here: covariations in the *latency* of response; and covariations in the *excitability* of response. Both are likely to be found in biological systems. This raises the possibility of covariograms which admit multiple, extremely different, interpretations: an interpretation problem which must be solved. The first step in solving it is to be aware of under what conditions interpretational ambiguity may arise (and, concomitantly, when it can be ruled out). Contributing to this has been the main objective of this paper. The second step is to resolve the ambiguity when it is present; some quantitative methods for doing so are proposed in the companion paper (Brody 1998; see also Friston 1995 and Vaadia et al. 1995). That excitability covariations could generate a peak in a JPSTH was a possibility raised by Aertsen et al. (1989), but they did not study the shape or magnitude of such a peak. Friston (1995; see also Vaadia et al. 1995) has described excitability covariations in more detail; similarities and differences between Friston’s work and that presented here are discussed in the companion paper (Brody 1998).

Three rules of thumb, for being alert to whether latency or excitability covariations could be present in a covariogram, may be gleaned from the examples and results of section 3:

Rule of thumb #1: Covariogram peak widths due to latency and excitability covariations are of the same order of magnitude as PSTH peak widths. This is due to the fact that excitability peaks are directly linked to terms containing stimulus-locked components⁶ (in addition to background firing

⁶If the variations in the gain factors balance out (that is, if $\langle \zeta^r \rangle = 0$), the PSTHs may be flat even in the presence of excitability covariations (Friston,

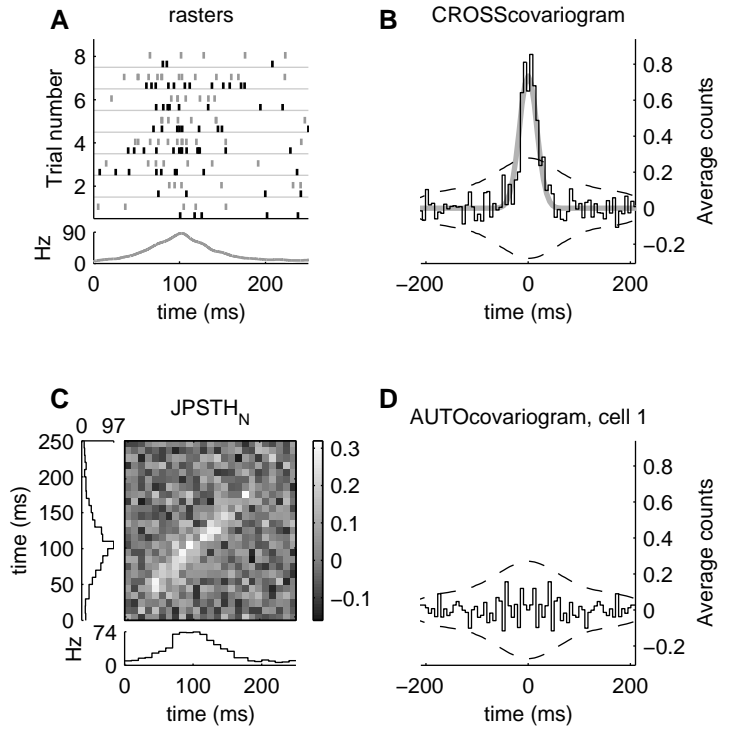


Figure 4: Spike timing covariations. **A:** Eight out of two hundred rasters used to illustrate spike timing covariations. Below the rasters are the smoothed PSTHs of the two cells. On each trial, the spike trains were made by first generating a single Poisson spike train, time-jittering the spikes of this twice, and then assigning the result of the first jittering to cell 1 and the result of the second jittering to cell 2. Additional spikes at a rate of 10 Hz were then added independently to each cell to represent background, uncorrelated, firing. [Details: The original spike train on each trial had firing rate $(70 \text{ Hz}) \exp(-(t - 100)^2 / (2 \cdot 30^2))$, with t in milliseconds. Jittering was done by adding a random amount of time to each spike, drawn from a Gaussian with mean 0 and standard deviation 12 ms.]. **B:** Covariogram of the two cells. The thick grey line is the analytical expected value of the covariogram, computed with knowledge of the parameters and procedures used to generate the spike trains. **C:** Normalised JPSTH of the two cells, flanked by their PSTHs. **D:** Autocovariogram of one of the cells; the other is similar. The central bin ($\tau = 0$) has been set to zero for display purposes. Unlike the latency and excitability cases (Figures 2 and 3), the autocovariogram does not resemble the crosscovariogram.

rate terms as wide as the entire trial itself— see equation 9). Latency peaks depend on R , the raw correlogram peak width, which in turn depends on the characteristic width of the cells’ responses (see Figure 2). Ken Britten (personal communication) has suggested estimating the characteristic width of the cells’ responses from the autocovariograms instead of the PSTHs. This leads to rule of thumb #2:

Rule of thumb #2: Latency and excitability covariations generate autocovariogram peaks that are similar to the crosscovariogram peaks. While spike timing covariations may exist without affecting the cells’ autocovariograms, latency or

1995).

excitability covariations add a contribution to the autocovariogram that is similar to their contribution to the crosscovariogram. This was shown here with Poisson-like, non-bursty model cells. The statement remains true if both cells are equally bursty. But if the cells are not symmetric (e.g. one is bursty but the other is not), the comparison between auto- and cross-covariograms will no longer be straightforward.

Rule of thumb #3: The integral (i.e. sum) of a covariogram is directly proportional to the covariation in the mean firing rates of the neurons (equation 12). Since the integral can be quickly estimated by eye, this measure should be in every covariogram—using neurophysiologist’s breast pocket.⁷ Large positive covariogram integrals indicate that the data was collected from trials with large positive covariations in their firing rates (implying the presence of an excitability covariations component), and suggests important changes of state during the experiment. Examples of covariograms with large positive integrals are common in the literature (e.g. Kruger and Aiple 1988; Alloway et al. 1993; Hata et al. 1993; Ghose et al. 1994; Sillito et al. 1994; Nowak et al. 1995; Munk et al. 1995). Note that spike timing covariations, as illustrated in Figure 4, can also generate positive covariogram integrals (common input can lead to spike timing coincidences, but can also lead to covariations in the *number* of spikes fired). However, in the spike timing case, the integral, if positive, will often be small, since the width of the peak can be very thin and unrelated to the width of the PSTH. Thus, the most tell-tale situation occurs when the integral is positive *and* the correlogram peak width is of the same order of magnitude as the PSTH peak widths. A straightforward method to determine whether the PSTHs can match the correlogram peak in this sense is presented in (Brody, 1998). (On the other hand, for examples of positive covariogram integrals clearly not caused by excitability covariations, see Tso et al. 1986.)

If any of the three rules of thumb suggest there could be important latency or excitability contributions to a covariogram, care should be taken before concluding that observed covariogram peaks are due to spike synchronization.

Further considerations

The three types of covariations that have been examined here are neither exhaustive nor exclusive. Other types of departure from independence also exist. Spike timing covariations may coexist with latency and/or excitability covariations, and latency and excitability covariations, in particular, may commonly exist in a paired manner: John Hopfield (personal communication) has suggested that covariations in resting membrane potential could lead to paired covariations in both latency and excitability, since depolarized resting potentials would lead to both high excitabilities and short latencies, while hyperpolarized resting potentials would lead to both low excitabilities and long latencies. Such changes in resting potentials might be induced by

variable ongoing activity in the network the neurons are part of (see Arieli et al. 1996).

All covariations were illustrated here using stochastic processes that were constant over all the trials of each simulated experiment. Differences between trials were simply different instantiations of the same stochastic process. Thus, there is no sense in which the process generating the spike trains for Figure 1A (spike timing) was any more, or less, stationary than those of Figure 1B or Figure 1C (latency and excitability). Nevertheless, in biological systems, variations in latency or excitability would most likely be due to slow changes of state, which are indeed nonstationarities. When Aertsen et al. (1989) mentioned interpretation problems associated with excitability covariations, they phrased them as due to nonstationarities.

As pointed out by an anonymous reviewer, the interpretation problems discussed in this paper may be seen as a special case of a more general problem, which is that of taking the mean of a distribution as representative of *all* the points of the distribution. Only when the standard deviation of a distribution is much smaller than its mean⁸ can the latter be meaningfully thought of as representative of the entire distribution; in biological systems, distributions are often broad and this condition is often not met. For example, the PSTH is defined as the average response over a set of trials, but if there are large variations in latency or excitability it is clearly not representative of each individual trial. Similarly, the covariogram is defined averaged over a set of trials, and should not necessarily be taken as representative of interactions occurring on each individual trial. Investigators must interpret means with care.

Acknowledgements

I am grateful to John Hopfield and members of his group, Sanjoy Mahajan, Sam Roweis and Erik Winfree, for discussion and critical readings of the manuscript. I also thank John Hopfield for support. I thank George Gerstein, Kyle Kirkland and Adam Sillito for discussion, and the anonymous reviewers for helpful comments. This work was supported by a Fulbright/CONACYT graduate fellowship and by NSF Cooperative Agreement EEC-9402726.

All simulations and analyses were done in MATLAB 5 (Mathworks Inc, Natick, MA). The code for all of these, including in particular the code to reproduce each of the figures, can be found at <http://www.cns.caltech.edu/~carlos/correlations.html>.

References

- A. M. H. J. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. Dynamics of neuronal firing correlation - modulation of effective connectivity. *J. Neurophysiol.*, 61(5):900–917, 1989.
- K. D. Alloway, M. J. Johnson, and M. B. Wallace. Thalamocortical interactions in the somatosensory system - interpretations of latency and cross-correlation analyses. *J. Neurophysiol.*, 70(3):892–908, 1993.
- A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen. Dynamics of ongoing activity - explanation of the large variability in evoked cortical responses. *Science*, 273(5283):1868–1871, 1996.
- C.D. Brody. Analysis and modeling of spike train correlations in the lateral geniculate nucleus. *Ph.D. Thesis, Caltech*, Available at <http://www.cns.caltech.edu/~carlos/thesis>, 1997a.

⁸For multivariate distributions, the square root of all eigenvalues of the covariance matrix must be much smaller than the magnitude of the mean.

⁷The integral of the covariogram is *exactly* proportional to the covariation in mean firing rates when $S_1^r(t)$ and $S_2^r(t)$ are defined to be zero for times outside trial r (section 2), for the purpose of computing the covariogram. If this is not done, the integral will include a term describing covariations in mean firing rates for times surrounding the trials. But even in this case, positive integrals should prompt investigators to look at covariations in mean rates.

- C.D. Brody. Latency, excitability, and spike timing correlations. *Society for Neuroscience Abstracts*, 23:14, 1997b.
- C.D. Brody. Disambiguating different covariation types. *Neural Computation*, **11**:1527-1535 (1999).
- M. Carandini and D. Ferster. A tonic hyperpolarization underlying contrast adaptation in cat visual-cortex. *Science*, 276(5314):949-952, 1997.
- K. J. Friston. Neuronal transients. *Proceedings of the Royal Soc. of London Series B Biological Sciences*, 261(1362):401-405, 1995.
- G. M. Ghose, I. Ohzawa, and R. D. Freeman. Receptive-field maps of correlated discharge between pairs of neurons in the cats visual-cortex. *J. Neurophysiol.*, 71(1):330-346, 1994.
- Y. Hata, T. Tsumoto, H. Sato, K. Hagihara, and H. Tamura. Development of local horizontal interactions in cat visual-cortex studied by cross-correlation analysis. *J. Neurophysiol.*, 69(1):40-56, 1993.
- J. Kruger and F. Aiple. Multimicroelectrode investigation of monkey striate cortex - spike train correlations in the infragranular layers. *J. Neurophysiol.*, 60(2):798-828, 1988.
- M. H. J. Munk, L. G. Nowak, J. I. Nelson, and J. Bullier. Structural basis of cortical synchronization .2. effects of cortical-lesions. *J. Neurophysiol.*, 74(6):2401-2414, 1995.
- L. G. Nowak, M. H. J. Munk, J. I. Nelson, A. C. James, and J. Bullier. Structural basis of cortical synchronization .1. 3 types of interhemispheric coupling. *J. Neurophysiol.*, 74(6):2379-2400, 1995.
- G. Palm, A. M. H. J. Aertsen, and G. L. Gerstein. On the significance of correlations among neuronal spike trains. *Biological Cybernetics*, 59(1): 1-11, 1988.
- D.H. Perkel, G.L. Gerstein, and G.P. Moore. Neuronal spike trains and stochastic point processes. ii. simultaneous spike trains. *Biophysics Journal*, 7:419-440, 1967.
- A. M. Sillito, H. E. Jones, G. L. Gerstein, and D. C. West. Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual-cortex. *Nature*, 369(6480):479-482, 1994.
- D. Y. Tso, C. D. Gilbert, and T. N. Wiesel. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.*, 6(4):1160-1170, 1986.
- E. Vaadia, A. Aertsen, and I. Nelken. 'dynamics of neuronal interactions' cannot be explained by 'neuronal transients'. *Proceedings of the Royal Soc. of London Series B Biological Sciences*, 261:407-410, 1995.