Instantaneous Cross-Correlation Analysis of Neural Ensembles with High Temporal Resolution[*]

António R. C. Paiva[1], Il Park[2], José C. Príncipe[3], and Justin C. Sanchez[4]

[1] Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT

arpaiva@sci.utah.edu

[2] Center for Perceptual System and Institute for Neuroscience, University of Texas, Austin, TX

memming@austin.utexas.edu

[3] Dept. of Electrical and Computer Engineering, University of Florida, Gainesville, FL

principe@cnel.ufl.edu

[4] Dept. of Biomedical Engineering, University of Miami, Coral Gables, FL

jcsanchez@miami.edu

Summary

One of the fundamental difficulties in neural assembly studies is the lack of an effective, high resolution measure of the spatio-temporal structure in spike trains obtained from a single realization. In this chapter a systematic approach to estimate the cross-correlation (CC) of spike trains, over time and in only one realization, is proposed. The solution lies in an alternate definition of cross-correlation which suggests that, rather than time averaging as is current practice, we should use ensemble averaging. This observation suggests a natural instantaneous CC estimator as required for high temporal resolution and real-time ensemble analysis and decoding. Results are shown in simulated datasets and neural activity of rat motor cortical neurons during a behavioral task.

1 Introduction

In multichannel neural studies, it is often necessary to quantify the spatio-temporal structure of spike train ensembles. How spike trains covary over time and at different scales is of high relevance to detect and analyze neural ensembles, study their functional connectivity, and their role in information encoding (Buzsáki, 2010; Freiwald et al., 2001; Strangman, 1997; Gerstein and

---

Turner, 1990; Gerstein et al., 1989). However, methods proposed in the literature to measure the covariation of spike trains neglect or smear heavily the evolution of this structure over time. This paper, however, proposes instantaneous cross-correlation as a high temporal resolution measure for multiple spike trains obtained from a single realization.

Correlation-based measures are widely used tools to analyze spike train covariation (Brown et al., 2004). In particular, the cross-correlation (CC) is often used to assess associations between neurons (Perkel et al., 1967; Brody, 1999). However, because of the nature of spike train data, the estimation of cross-correlation is not straightforward due to the lack of the concept of signal amplitude. To deal with this problem, CC is commonly computed on the binned spike trains. In effect, binning is a transformation of the uncertainty in time onto amplitude variability and creates a form of averaging over time. In addition, binning limits the time resolution to the bin size due to the implicit quantization of the spike times. Moreover, to improve the estimation, local stationarity and ergodicity are tacitly assumed to allow averaging over a window in time. Although this averaging improves the estimation accuracy, this is obtained at the expense of temporal resolution since the filtering smears the temporal dynamic of the coupling.

Many other methods for analysis of spike train covariation have been proposed in the literature, including the JPSTH (Gerstein and Perkel, 1972; Aertsen et al., 1989), unitary events (Grün et al., 2002a,b), and frequency methods (Baccalá and Sameshima, 1999; Sameshima and Baccalá, 1999), among others. In any of these methods, however, the time structure of the analysis is smeared by averaging, over time or over trials, for gathering of statistics, of through the transformation involved.

To obtain an estimator of spike train cross-correlation with high temporal resolution, we must avoid averaging over time. The solution to this problem lies in the definition of a spatial cross-correlation. The crucial observation is that, rather than time averaging as is commonly done using the binned spike train CC estimator, we should use ensemble averaging. This observation suggests a natural instantaneous CC estimator as required for high temporal resolution and real-time ensemble analysis and decoding.

For spike train analysis we are particularly interested in measuring temporal structure in the coupling between neurons. Inevitably, this is often associated with the formation of neural assemblies and for which the most widely used descriptor is synchrony in the neuron firings (Gerstein and Turner, 1990; Gerstein et al., 1989). Due to the importance of precise spike times, synchrony studies in neural assemblies arise naturally as an obvious application for the ICC. Yet, the use of the ICC is not limited to this situation, as will be clear in our presentation. By selecting the analysis parameter appropriately it is possible to control the time scale of analysis covering the full spectrum, from synchronous firings to covariation in firing rate modulation.

2   Background

The cross-correlation (CC) function of spike trains is typically computed on the binned spike trains (Brown et al., 2001). The practical reasons for this approach are easily understood. By binning the spike trains one obtains a discrete-time random process, and thus random process theory can be applied as usual. Hence, the CC function of binned spike trains is defined at lag l as

$$C_{AB}^{bin}[l] = E\left\{b_A[n]b_B[n+l]\right\}, \tag{1}$$

where b_A, b_B denote the binned spike trains, with $b_A[n]$ and $b_B[n]$ corresponding to the number of spikes in the nth bin for spike train A and B, respectively. Additionally, in practice, the expectation is approximated as an average over time (i.e., bins),

$$C_{AB}^{bin}[l] \approx \frac{1}{M}\sum_{n=1}^{M} b_A[n]b_B[n+l], \tag{2}$$

with M the number of bins. To observe the time evolution of the cross-correlation, the averaging is usually done using a sliding window in time.

As stated earlier, a fundamental problem with the use of this approach for the analysis of temporal structure is that the averaging over time will hide any existing temporal coupling dynamics at the fine time scale. Statistically, averaging over time implies that ergodicity and stationarity are assumed, at least locally. Yet, these assumptions are contrary to the goal of estimating the temporal dynamics.

Another issue with the above definitions is the use of binned spike trains. It is important to note that binning of spike trains imposes a quantization on the spike times. Put differently, any information that is expressed at higher resolutions (i.e., finer time scale) than the bin size will be lost through binning. This is especially concerning when one is interested in analysis of fine temporal structure, in the form of synchrony for example (Grün et al., 1999).

3   Instantaneous Cross-Correlation

3.1   Preserving temporal resolution

The goal here is to derive a CC estimator with high temporal resolution. We formalize this requirement in the CC definition in equation (1) by explicitly showing the dependence on the time instant n, yielding the instantaneous cross-correlation (ICC),

$$C_{ij}^{bin}[n,l] = E\left\{b_i[n]b_j[n+l]\right\}. \tag{3}$$

Written in this way, the above definition reinforces the idea that the expectation must be estimated over realizations rather than over time in order to preserve the time dynamics.

Nevertheless, one must still decide on how to define a realization to approximate the expectation. In a multi-trial experimental paradigm one could approximate the expectation by averaging over trials, as done using the PSTH (Gerstein and Kiang, 1960) and JPSTH (Gerstein and Perkel, 1972; Aertsen et al., 1989). Averaging over trials preserves the trial temporal dynamics *only if the trials were stationary*. However, this assumption is hard to verify in practice and the analysis' results can deviate significantly if the assumption does not hold (Brody, 1999). In a single-trial paradigm, and to avoid the limiting assumptions of stationarity in general, here it is proposed to approximate the expectation by the mean over all pairs of spike trains. Consider spike trains $\{s_1, \ldots, s_N\}$, then the ICC can be estimated as

$$\hat{C}^{bin}[n, l] = \frac{2}{N(N-1)} \sum_{i=1}^{N} \sum_{j=i+1}^{N} b_i[n] b_j[n+l]. \tag{4}$$

Note that, in a sense, this approach is actually statistically more efficient than time averaging because the number of spike train pairs increases quadratically with the number of spike trains rather than linearly.

The ensemble averaged ICC is a spatio-temporal measure of the ensemble covariation over time. In this form, and due to the exchange of time for ensemble averaging, the ICC is capable of detecting the presence of temporal dynamics within the ensemble with high temporal resolution. However, it raises the problem of how to select the subset of averaging neurons comprising an ensemble. A potential limitation in the use of the ICC is that spatial and neuron specificity is lost in the process of ensemble averaging. This is the trade-off incurred to reduce the estimation noise and still preserve the high temporal resolution.

3.2 Avoiding time quantization

The previous section introduced the ICC as an estimator of ensemble cross-correlation over time, but it did not address the issue of time quantization imposed by using binning. As stated earlier, any information expressed in the spike times at a finer time scale than the bin size will be lost through binning. Moreover, as pointed out by Richmond et al. (1990), the binning boundaries create an arbitrary assignment error. One might think that the error associated with binning is negligible, but recall that time quantization limits the resolution and significance of the analysis, which is precisely the goal here.

One should then question, what is binning actually doing? And, correspondingly, is there a better way to do it? As noted by Dayan and Abbott (2001), binning is an estimator of the

instantaneous firing rate (apart from normalization by the bin size). Hence, the cross-correlation between two spike trains in equation (3) can be written in the more general form

$$C_{ij}(t,\theta) = E\left\{\hat{\lambda}_i(t)\hat{\lambda}_j(t+\theta)\right\} \tag{5}$$

where $\hat{\lambda}_i(t)$ and $\hat{\lambda}_j(t)$ denote the intensity functions estimated from spike trains i and j. This definition has the advantage that any intensity function estimator can be used. Specifically, in the statistical literature the conventional approach for intensity function estimation of point processes is kernel smoothing Reiss (1993), with clear advantages in the estimation Kass et al. (2003); Richmond et al. (1990). Denoting the spike times in the time interval $[0, T]$ of the ith spike train as $\{t_n^i : n = 1, \ldots, N_i\}$, where N_i is the number of spikes of A in the interval. The kernel estimated intensity function is given by

$$\hat{\lambda}_i(t) = \sum_{n=1}^{N_i} h(t - t_n^i), \tag{6}$$

where h is the smoothing kernel function, which must integrate to one. Note that, unlike binning, kernel intensity estimation does not introduce time quantization errors.

Equation (6) can be interpreted as the convolution of the spike train with a window given by the smoothing function h. By controlling the smoothing introduced by h, one can regulate the scale at which the intensity function is estimated. Therefore, this controls the time-scale of the CC coupling between two extremes: synchrony in neuron firings (for small smoothing) or firing rate (for large smoothing).

Nevertheless, it must be emphasized that the smoothing in equation (6) is fundamentally different from the time averaging in the estimation of CC. The smoothing in the estimation of the intensity function defines the time-scale at which the spike times are analyzed for any information they may express. The CC time averaging, on the other hand, is used to reduce the estimation noise in order to improve the statistical significance of the result. Of course, if the smoothing introduced by h is large, there is a reduced need for time averaging, but this is simply because changes in firing rate must inherently vary slower than changes in synchrony.

Although many different smoothing functions h can be utilized, estimation methodologies that are causal and can be applied online are of particular interest. A smoothing function with these properties is the exponential function,

$$h(t) = (1/\tau)\exp\left(-t/\tau\right)u(t), \tag{7}$$

where $u(\cdot)$ is the step function. The exponential function provides both graded interactions and a time scale for the intensity estimation by controlling the time constant τ. In addition, the

exponential function is biologically plausible since it can be interpreted as the evoked post-synaptic potential in a neuron Dayan and Abbott (2001). Accordingly, the estimated intensity function for the ith spike train is given by

$$\hat{\lambda}_i(t) = (1/\tau) \sum_{t_n^i \leq t} \exp\left[-(t - t_n^i)/\tau\right] u(t - t_n^i),$$ (8)

which is simply the filtering of a spike train by a first order IIR filter and, therefore, it is computational very simple and particularly suitable for hardware implementation.

Using the kernel intensity function estimator of equation (6), the ICC is defined as

$$\hat{C}(t, \theta) = \frac{2}{N(N-1)} \sum_{i=1}^{N} \sum_{j=i+1}^{N} \hat{\lambda}_i(t)\hat{\lambda}_j(t + \theta).$$ (9)

Hereafter, ICC refers to this formulation unless explicitly stated otherwise.

4   Experiments

4.1   Synchronization of pulse-coupled oscillators

In this example, we show that ICC can quantify the increase in coupling strength in a spiking neural network of leaky-integrate-and-fire (LIF) neurons designed according to Mirollo and Strogatz Mirollo and Strogatz (1990)[1] and the ICC results compare favorably with the extended cross-correlation for multiple neurons. The network is initialized in a random condition and is proven to synchronize over time. The synchronization is essentially due to leakiness and the weak coupling among the oscillatory neurons.

The raster plot of neuron firings is shown in Fig. 1. There are two main observations: the progressive synchronization of the firings associated with the global oscillatory behavior of the network, and the local grouping that tends to preserve local synchronizations that either entrain the full network or wash out over time, as expected from theoretical studies of the network behavior Mirollo and Strogatz (1990). The ICC depicts this behavior precisely: the synchronization increases monotonically, with a period of fast increase in the first second followed by a plateau and slower increase as time advances. Moreover, it is possible to observe in the first 1.5s the formation of a second group of synchronized neurons which slowly merges into the main group. In addition, the envelope of ICC reveals the coherence of the membrane potentials as quantified by the information potential (IP). The IP is an information theoretic quantity inversely proportional

---

[1]The parameters for the simulation are: 100 neurons, resting and reset membrane potential -60mV, threshold -45mV, membrane capacitance 300nF, membrane resistance 1MΩ, current injection 50nA, synaptic weight 100nV, synaptic time constant 0.1ms and all to all excitatory connection.
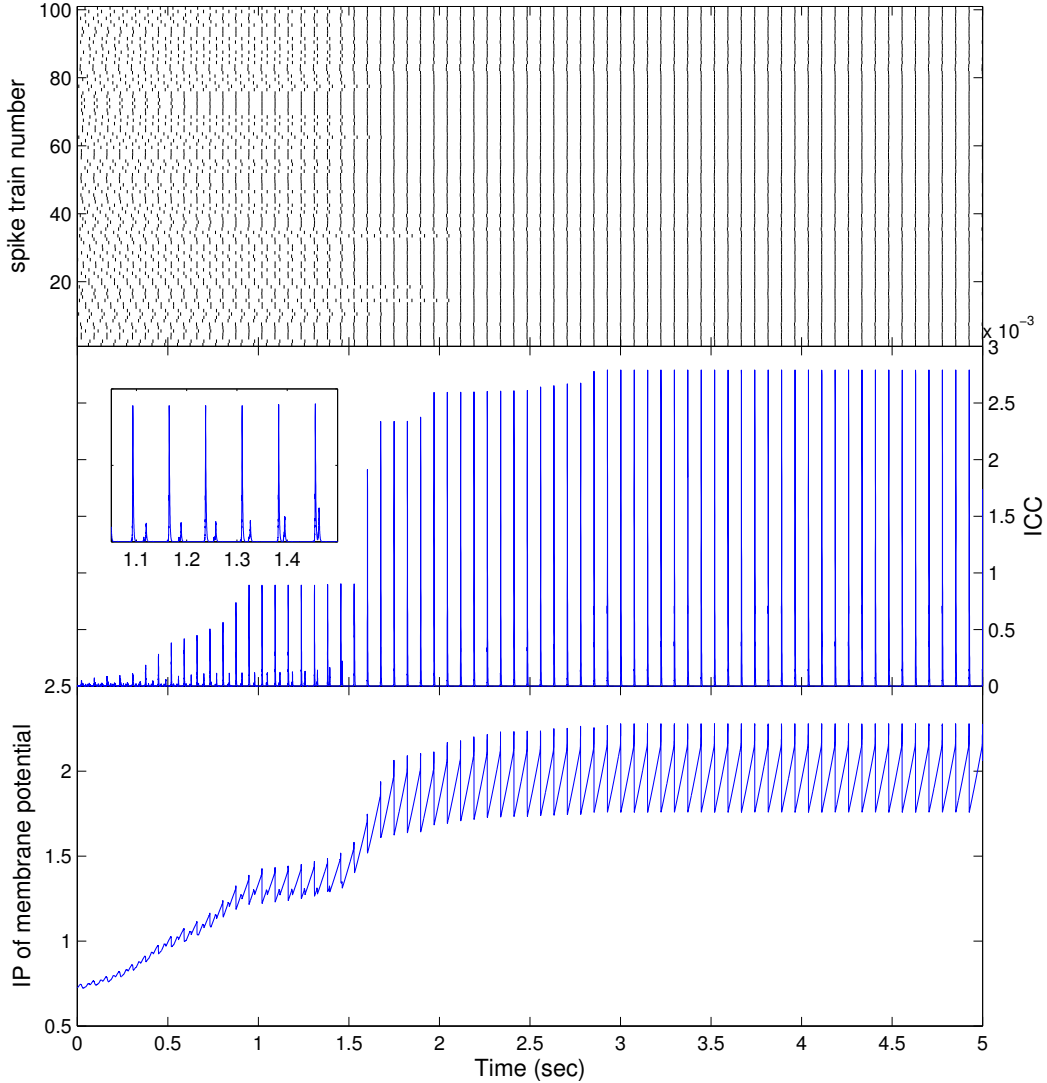
Figure 1: Evolution of synchrony in a spiking neural network. (Top) Raster plot of the neuron firings. (Middle) ICC over time. The inset highlights the merging of two synchronous groups. (Bottom) Information potential of the membrane potentials. This is a macroscopic variable characterizing the entropy of the neurons' internal state.

to entropy Principe (2010). It was computed with $\mathrm{IP}_\theta = \frac{1}{M^2} \sum_{i=1}^{M} \sum_{j=1}^{M} \exp(-d(\theta_i, \theta_j)/2\sigma^2)$ with $\sigma = 75\mathrm{mV}$.[2] The IP measures the coherence of the neuron's internal state, which is only available in simulated networks. Yet, it is interesting to verify that the ICC was able to successfully and accurately characterize such information from the observed spike trains.

For comparison, the zero-lag cross-correlation over time, averaged through all pairwise combinations of neurons, is presented in Fig. 2. The cross-correlation was computed with a sliding

---

[2]The distance used in the Gaussian kernel was $d(\theta_i, \theta_j) = \min\left(|\theta_i - \theta_j|, 15mV - |\theta_i - \theta_j|\right)$, where θ_i is the membrane potential of the ith neuron. This wrap-around effect expresses the phase proximity of the neurons before and after firing.
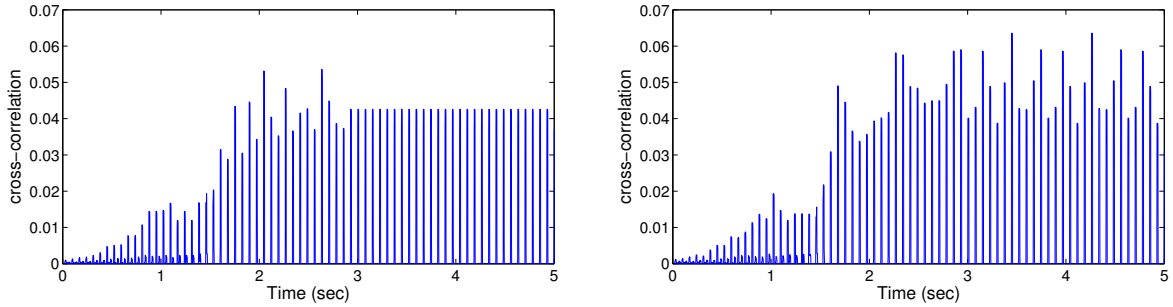
Figure 2: Zero-lag cross-correlation computed over time using a sliding window 10 bins long, and bin size 1ms (left) and 1.1ms (right).

window 10 bins long, sliding 1 bin at a time. Results are shown for bin sizes of 1ms and 1.1ms. It is noteworthy that although cross-correlation captures the general trends of synchrony, it masks the plateau and the final synchrony, and it is highly sensitive to the bin size as highlighted by the figure. These results highlight the importance of working in continuous-time, as proposed in Section 3.2, in contrast to the use of binning.

4.2 Analysis of neural synchronous activity in motor neurons

In our second example, the ICC is utilized to analyze the presence of synchronous activity in the motor cortex of a rat's brain. Throughout the literature, synchronous activity has been shown to provide additional information about motor movement when compared to firing rate modulation pattern analysis alone, including when no firing rate modulations are noticeable (Vaadia et al., 1995; Hatsopoulos et al., 1998; Riehle et al., 1997). Indeed, synchronous neural activity seems to be an widespread characteristic of the brain and can be found in a number of cortices, such as the auditory (Wagner et al., 2005; Carr and Konishi, 1990) and the visual cortices (Freiwald et al., 2001), for example.

Multichannel neuronal firing times from a male Sprague-Dawley rat were simultaneously recorded during a conditioned behavioral task at the University of Florida McKnight Brain Institute. The rat was chronically implanted with two 2×8 arrays of micro-electrodes placed bilaterally in the forelimb region of the primary motor cortex (1.0mm anterior, 2.5mm lateral of bregma (Donoghue and Wise, 1982)). Neuronal activity was collected with a Tucker-Davis recording rig with sampling frequency of 24414.1Hz and digitized to 16 bits of resolution. The firing times were recorded from individual neurons spike sorted with an online algorithm employing a combination of thresholding and template-based techniques. From sorting, a total of 44 single neurons were recorded, 24 neurons from the left hemisphere and 20 neurons from the right hemisphere. Simultaneously, the rat performed a go no-go lever pressing task in an operant conditioning cage (Med-Associates,
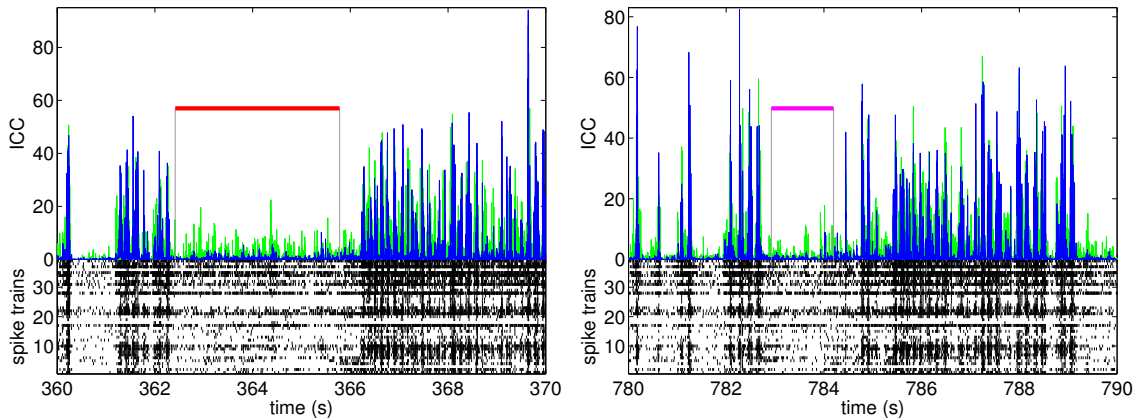
Figure 3: ICC and neuron firing raster plot on a single realization, showing the modulation of synchrony around the lever presses. The ICC was averaged throughout all neurons pairs separately for each hemisphere: left (blue) and right (green). The left plot shows a left lever press and the right plot shows a right lever press.

St. Albans, VT, USA). The task consisted of choosing and pressing one out of two levers (left or right) depending on a LED visual stimulus to obtain a water reward. The queue and lever press signals were recorded simultaneously with the neural activity with sampling frequency 381.5Hz. See Sanchez et al. (2005) for additional details on the experimental configuration.

ICC was applied to this dataset to quantify the synchronous neural activity across the ensemble. Figure 3 shows two trials with the ICC at zero-lag. As verified in other trials, the figure shows an increase in synchrony when the lever is released. Moreover, while the lever was being pressed the ensemble synchrony was observed to be significantly smaller than synchrony before and after the lever press. Actually, in this dataset, visual inspection of the raster plots would yield such conclusions, but ICC provides a quantitative method to translate the visual evaluation. Furthermore, examining the raster plots we can verify the presence of ensemble synchronized activity reoccurring in a periodic manner after the lever is released. Notice that the ensemble ICC captures the presence of this oscillatory synchronized activity, seen in the envelope of ICC in Figure 3, directly from a single trial and with high temporal resolution.

To investigate either the observed synchrony modulation was widespread across all the neurons, we clustered the neural activity after the lever release was clustered. For that purpose, we used the spike train clustering algorithm described in Paiva (2008); Paiva et al. (2009) with the memoryless cross-intensity (mCI) kernel. Figure 4 shows that all neurons contribute to the observed synchrony modulation, even though different clusters have slightly different rates and dynamics. In other words, neurons tend to form overlapping and time-evolving assemblies (Buzsáki, 2010; Strangman, 1997). Generally, clustering could be applied first to ensure that the ICC spatially averages only within an ensemble and, in this example, it could help to quantify sub-assembly dynamics.
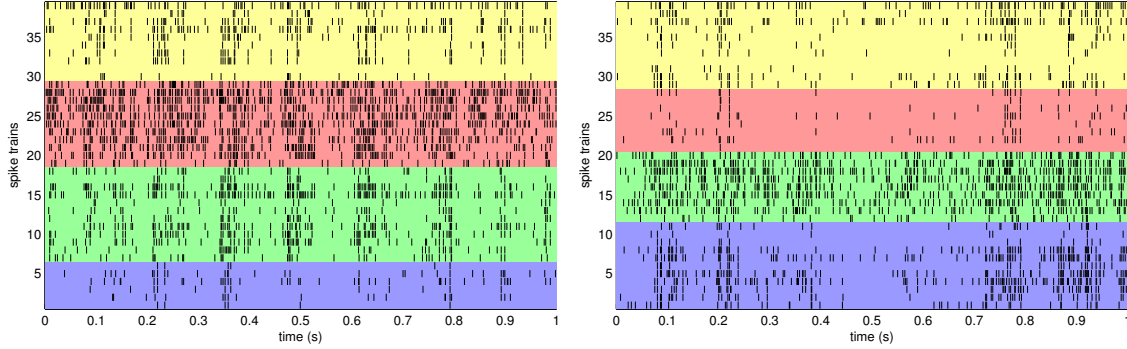
Figure 4: Clustering of neural activity following the lever release, considering 4 clusters, for the same trials shown in Figure 3.
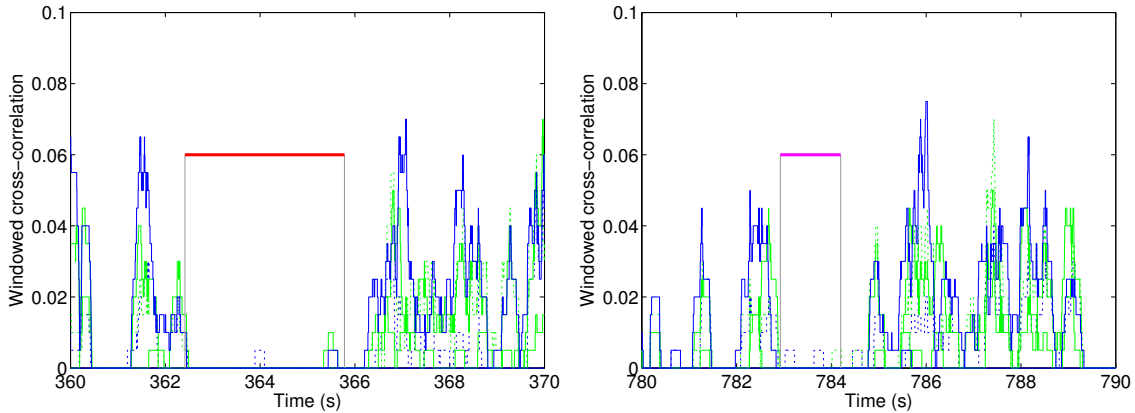

Figure 5: Windowed cross-correlation of selected 6 pairs of neurons, for the same segments shown in Figure 3. The cross-correlation was computed with a 200ms sliding window over 1ms bins. Four of the neuron pairs, two from each hemisphere, are known to synchronize and are shown with a solid line. The remaining two, one from each hemisphere, do not synchronize strongly and are shown with a dotted line.

For comparison with the commonly used methodologies, we show the cross-correlation computed at zero-lag with a 200ms sliding window over 1ms bins, as is commonly done in other studies (see, for example, Hatsopoulos et al. (1998)). This is shown first for only some selected pairs of neurons (Figure 5), and then spatially averaged (Figure 6) as proposed in Section 3.1. Although the presence of synchrony is also successfully captured with cross-correlation, the existence of a periodic modulation in synchrony is not noticeable. This can be expected since the cross-correlation requires stationarity over time and uses time averaging to reduce the randomness of the estimator. These two factors filter out potentially interesting features, such as the modulation of synchrony, on which information may represented. Put differently, time averaging imposes up front a lower bound on the frequencies that can analyzed. This effect is most visible in Figure 6 where spatial averaging greatly improves the estimation as the variance of the estimator is reduced, but the time
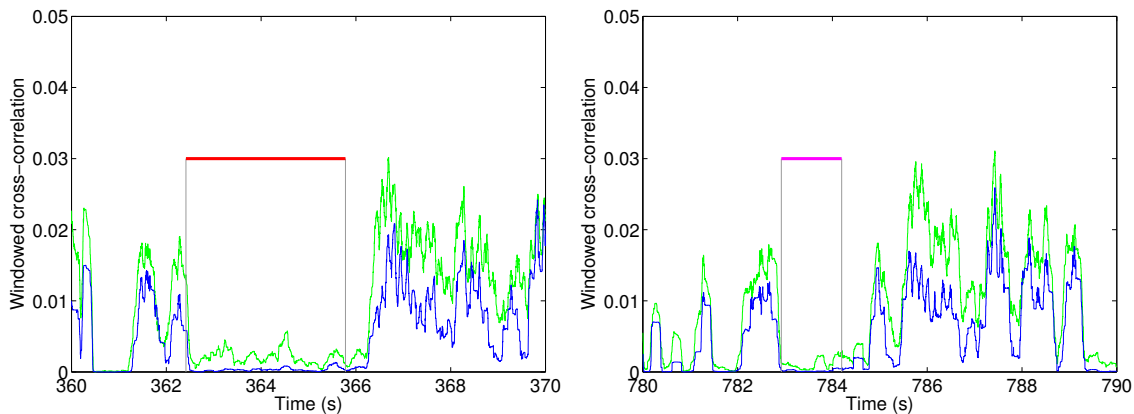
10

Figure 6: Ensemble averaged windowed cross-correlation, for the same segments shown in Figure 3. The cross-correlation for each neuron pair was computed with a 200ms sliding window over 1ms bins. As proposed in Section 3.1, the spatial average was taken throughout all neuron pair combinations, separately for each hemisphere: left (dark gray) and right (light gray).

averaging masks the modulation in synchrony in favor of general trends. These figures highlight the importance of the spatial averaging proposed for ICC, in opposition to the time averaging employed in cross-correlation.

5 Discussion

This chapter aims primarily (i) to highlight the limitations of cross-correlation as it is commonly estimated and (ii) to introduce the ICC as an alternative measure of cross-correlation with high temporal resolution. As reviewed in Section 2, the current practice for CC estimation utilizes time averaging and binning. Yet, this approach limits the temporal resolution of the analysis and the dynamics which can be characterized by the measure (cf. Section 4.1). Perhaps most importantly, the binned time-averaged CC estimator requires strong assumptions about the stationarity and ergodicity of the neural activity. However, if these assumptions do not hold, the outcome of the data analysis will likely have little or no significance (Brody, 1999). In contrast, the ICC uses ensemble averaging to obtain a measure of the neurons' cross-correlation with high temporal resolution. Using ensemble averaging has the added advantage that the approximation to the expectation has generally smaller variance because the number of spike train pairs increases quadratically. Thus, ICC can directly take full advantage of the increase in the channels in multi-channel spike train recordings. In addition, binning artifacts are avoided by computing the ICC in continuous time, as defined by equation (9) and demonstrated in the experiments.

On the other hand, using ICC involves a change in analysis perspective because spatial and neuron specificity is lost as a result of ensemble averaging. However, this is a necessary trade-off

because, as stated by the Heisenberg uncertainty principle, an increase in time resolution implies a decrease in spatial resolution. This is not a limitation, however, if the goal is to analyze the dynamic coupling behavior of the neural ensemble Gerstein and Turner (1990). In any case, it raises the question on how to divide the spike train recordings in neural ensembles. Although further tests still need to performed, the solution might involve clustering the spike trains (Paiva et al., 2009; Paiva, 2008), as illustrated in Section 4.2.

The ICC is presented here as a data analysis tool to the practitioner, but it should be pointed out that the ICC also provides the framework from computation and adaptation in the spike train domain. Indeed, the ICC defines an inner product between spike trains at different times. This follows from the fact that, by definition, the ICC is a correlation function and therefore it is an inner product (Parzen, 1959). More explicitly, one can verify that the ICC is a symmetric and positive definite function and thus verifies the conditions of the Moore-Aronszajn theorem (Aronszajn, 1950). In other words, since the ICC is symmetric and positive definite, it induces a reproducing kernel Hilbert space (RKHS) for which the ICC evaluates the inner product. Having the mathematical structure of an RKHS means that there exists a complete vector space on which we can compute with spike trains directly and apply PCA, clustering, Fisher discriminant analysis, etc. (Paiva et al., 2009; Paiva, 2008; Paiva et al., 2010). The inner product defined by the ICC is particularly interesting because it combines the pairwise cross-correlations to characterize the spike trains' joint distribution. Compared to existing statistical methods (Truccolo et al., 2005), this is a significant achievement considering that this information can be extracted directly, without knowledge of the data model, and in a single trial.

References

A. M. 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.

N. Aronszajn. Theory of reproducing kernels. *Trans. Am. Math. Soc.*, 68(3):337–404, May 1950.

Luiz Antonio Baccalá and Koichi Sameshima. Directed coherence: a tool for exploring functional interactions among brain structures. In Miguel A. L. Nicolelis, editor, *Methods for Neural Ensemble Recordings*, pages 179–192. CRC Press, 1999.

Carlos D. Brody. Correlations without synchrony. *Neural Comp.*, 11(7):1537–1551, 1999.

E. N. Brown, D. P. Nguyen, L. M. Frank, M. A. Wilson, and V. Solo. An analysis of neural receptive field plasticity by point process adaptive filtering. *Proc. Natl. Acad. Sci.*, 98:12261–12266, 2001.

Emery N. Brown, Robert E. Kass, and Partha P. Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neurosci.*, 7:456–461, 2004. doi: 10.1038/nn1228.

G. Buzsáki. Neural syntax: cell assemblies, synapsembles, and reader. *Neuron*, 68:362–385, November 2010. doi: 10.1016/j.neuron.2010.09.023.

C. E. Carr and M. Konishi. A circuit for detection of interaural time differences in the brain stem of the barn owl. *J. Neurosci.*, 10(10):3227–3246, 1990.

Peter Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, MA, USA, 2001.

J. P. Donoghue and S. P. Wise. The motor cortex of the rat: cytoarchitecture and microstimulation mapping. *J. Comp. Neurol.*, 212(12):76–88, 1982.

Winrich A. Freiwald, Andreas K. Kreiter, and Wolf Singer. Synchronization and assembly formation in the visual cortex. In M. A. L. Nicolelis, editor, *Progress in Brain Research*, pages 111–140. Elsevier Science, 2001.

G. L. Gerstein and N. Y.-S. Kiang. An approach to the quantitative analysis of electrophysiological data from single neurons. *Biophys. J.*, 1(1):15–28, September 1960.

G. L. Gerstein and D. H. Perkel. Mutual temporal relationships among neuronal spike trains. statistical techniques for display and analysis. *Biophys. J.*, 12(5):453–473, May 1972.

George L. Gerstein and Mark R. Turner. Neural assemblies as building blocks of cortical computation. In Eric L. Schwartz, editor, *Computational Neuroscience*, pages 179–191. MIT Press, 1990.

George L. Gerstein, Purvis Bedenbaugh, and Ad M. H. J. Aertsen. Neuronal assemblies. *IEEE Trans. Biomed. Eng.*, 36(1):4–14, January 1989.

Sonja Grün, Markus Diesmann, Franck Grammont, Alexa Riehle, and Ad Aertsen. Detecting unitary events without discretization of time. *J. Neurosci. Methods*, 93(1):67–79, December 1999. doi: 10.1016/S0165-0270(99)00126-0.

Sonja Grün, Markus Diesmann, and Ad Aertsen. Unitary Events in multiple single-neuron activity. I. detection and significance. *Neural Comp.*, 14(1):43–80, 2002a.

Sonja Grün, Markus Diesmann, and Ad Aertsen. Unitary Events in multiple single-neuron activity. II. nonstationary data. *Neural Comp.*, 14(1):81–119, 2002b.

N. G. Hatsopoulos, C. L. Ojakangas, L Paninski, and J. P. Donoghue. Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proc. Natl. Acad. Sci.*, 95(26):15706–15711, December 1998.

Robert E. Kass, Valérie Ventura, and Can Cai. Statistical smoothing of neuronal data. *Netw.: Comp. Neural Syst.*, 14:5–15, 2003.

Renato E. Mirollo and Steven H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM J. Applied Math.*, 50(6):1645–1662, December 1990.

António R. C. Paiva. *Reproducing Kernel Hilbert Spaces for Point Processes, with Applications to Neural Activity Analysis*. PhD thesis, University of Florida, 2008.

António R. C. Paiva, Il Park, and José C. Príncipe. A reproducing kernel Hilbert space framework for spike train signal processing. *Neural Comp.*, 21(2):424–449, February 2009. doi: 10.1162/neco.2008.09-07-614.

António R. C. Paiva, Il Park, and José C. Príncipe. Inner products for representation and learning in the spike train domain. In Karim G. Oweiss, editor, *Statistical Signal Processing for Neuroscience and Neurotechnology*, chapter 8. Academic Press, 2010.

Emanuel Parzen. Statistical inference on time series by Hilbert space methods. Technical Report 23, Applied Mathematics and Statistics Laboratory, Stanford University, Stanford, California, January 1959.

D. H. Perkel, G. L. Gerstein, and G. P. Moore. Neuronal spike trains and stochastic point processes. II. simultaneous spike trains. *Biophys. J.*, 7(4):419–440, July 1967.

Jose C. Principe. *Information Theoretic Learning.* Springer, 2010. ISBN 978-1-4419-1569-6.

Rolf-Dieter Reiss. *A Course on Point Processes.* Springer-Verlag, New York, NY, 1993.

Barry J. Richmond, Lance M. Optican, and Hedva Spitzer. Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus-response relations. *J. Neurophysiol.*, 64(2):351–369, August 1990.

Alexa Riehle, Sonja Grün, Markus Diesmann, and Ad Aertsen. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278(5345):1950–1953, December 1997. doi: 10.1126/science.278.5345.1950.

Koichi Sameshima and Luiz Antonio Baccalá. Using partial directed coherence to describe neuronal ensemble interactions. *J. Neurosci. Methods*, 94(1):93–103, December 1999.

J. C. Sanchez, J C. Principe, and P. R. Carney. Is neuron discrimination preprocessing necessary for linear and nonlinear brain machine interface models? In *Proc. Intl. Conf. on Human-Computer Interaction*, 2005.

Gary Strangman. Detecting synchronous cell assemblies with limited data and overlapping assemblies. *Neural Comp.*, 9(1):51–76, 1997.

Wilson Truccolo, Uri T. Eden, Matthew R. Fellows, John P. Donoghue, and Emery N. Brown. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.*, 93:1074–1089, February 2005. doi: 10.1152/jn.00697.2004.

E. Vaadia, I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A. Aertsen. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373(6514):515–518, February 1995. doi: 10.1038/373515a0.

Hermann Wagner, Sandra Brill, Richard Kempter, and Catherine E. Carr. Microsecond precision of phase delay in the auditory system of the barn owl. *J. Neurophysiol.*, 94(2):1655–1658, 2005. doi: 10.1152/jn.01226.2004.