Asymmetric sliding-window cross-correlation

Alex Bäcker and Stijn Cassenaer

Introduction

Sliding-window cross-correlation is a common method to esimate time-varying correlations between signals (Laurent and Davidowitz, 1994; Laurent et al., 1996; Macleod and Laurent, 1996; Stopfer and Laurent, 1997; Wehr, 1999 (p. 96)). It produces a correlation value betwen two signals (positive or negative) for every (time,lag) pair of values. In principle, the expected value of the correlation for any pair of (time,lag) values must be computed by averaging x(t).y(t+lag) over many realizations of the stochastic process. This is called an ``ensemble average'' across realizations of a stochastic process. Lacking a large enough number of realization over which to average, one must resort to other methods. If the correlations are stationary (i.e. time-invariant), then we may average across time to estimate the expected value. If the correlations are not stationary, we may divide the signals into sliding windows of a size such that the correlations can be considered stationary on the timescale of the window width, and calculate the cross-correlation, as a function of lag, for each window, sliding the window along the signal to obtain correlations for different time values. This is what is called a sliding-window cross-correlogram.

The problem

Take, for example, two signals which start off uncorrelated, then show a periodic positive correlation for 1 second, and finally become uncorrelated again. The cross correlogram at the beginning of the 1-second period of correlation should show positive correlation with positive lags, but not with negative lags: the signal is correlated to what will come, but not to what was there before. In the middle of the 1-second period, correlation is positive for negative and positive lags. Toward the end, correlation is positive only for negative lags. And yet none of the sliding-window cross-correlograms in the literature exhibit this asymmetry.

The reason why the asymmetry is lacking is the following: the correlograms were computed by calculating the cross-correlation function for each window separately, using commercial routines such as MATLAB's *xcorr* function, which slide one signal's vector past the other one for each window. With this method, only the values within the window being used are used for the correlation, and thus whether the window is before, after or in the middle of a period with high correlation makes no difference other than by the correlation present in the window itself. In particular, if two signals, s and s', are perfectly correlated, as in an autocorrelogram, their cross-correlation will be symmetrical for each window (i.e. for each t-value, taken to lie in the middle of each window) by construction:

c(t,+lag) = s(t-lag/2).s'(t+lag/2)

c(t,-lag) = s(t+lag/2).s'(t-lag/2) = s'(t-lag/2).x(t+lag/2)

If s is perfectly correlated with s', c(t,+lag)=c(t,-lag), because the asymmetry will show up in one signal at positive lag and in the other signal at negative lag*. In other words, instead of having the correlation at positive lag computed from a comparison with a window shifted in the positive direction and the correlation at negative lag computed from a comparison with a window shifted in the negative direction, the existing method uses the same window for both lags, simply shifting the window in different directions.

That method leads to another, related, problem: the greater the magnitude of the lag, the less data is used. This happens because the edge of the window does not move as the lag is changed, and thus only lag zero allows a comparison between every sample in the window for each signal. For lags of any magnitude, the shift between signals forces the comparison to be done over every decreasing stretches of signal, until at lags of the window length, a single sample from each signal is used. Thus, the traditional method will yield noisy correlation estimates for any lags which are not significantly less than the length of the windows used (see Fig. 1), and cannot be used at all for lags greater than the window length. The most powerful cross-correlogram, though, is one which has small window length (so as not to blurr variations in time) and large lag ranges (to observe correlations at any lag). In particular, the windows have to be small compared to the timecourse of variations in the correlation. This means that at the onset and offset of the oscillations, the windows should be particularly small. But if the maximum lag is constrained to the length of the window, small windows do not allow seeing the periodic structure of the correlation.

A solution

As discussed above, the motivation behind sliding windows is the assumption that signals are relatively stationary on the timescale of the window length. Because there are no hard boundaries, the signal is not stationary only within the windows, but rather on the timescale of the window length, and thus samples toward the edge of the window should be correlated with samples of the other signal within a window-length from them. In other words, if one knows the signal surrounding a window, those samples must be taken into account in calculating the mean signal following or preceding samples in the window. Thus, the cross correlation between s and s' at time t is given by:

c(t,lag) = < s(t).s'(t+lag) >,

where the average is over all t values in the window.

This method, which we term asymmetric cross-correlation, has several advantages over its predecessor, termed symmetric cross-correlation below for comparison: 1) it eliminates the artificial symmetry, 2) it eliminates the reduction of data for increased lag magnitudes, and 3) it allows for small windows concurrently with large lag ranges[1]. This is illustrated in figure 2.

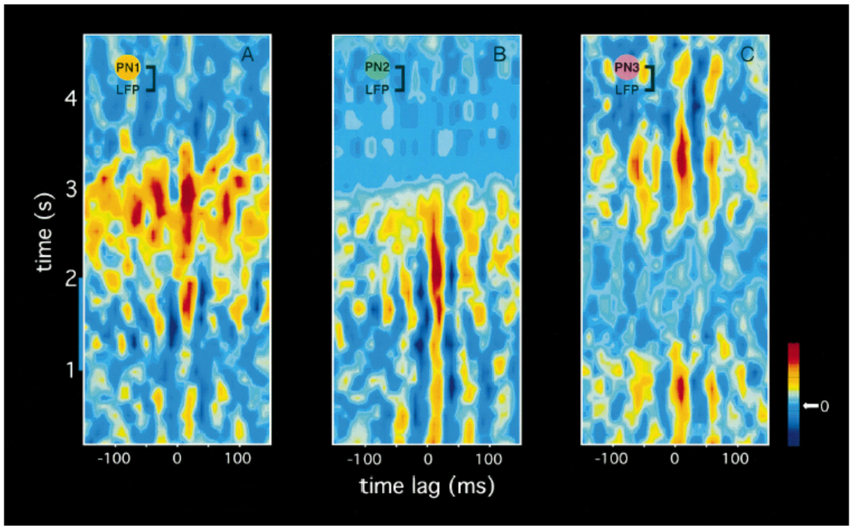

Figure 1. Sliding-window cross-correlograms (from Laurent et al., 1996) calculated with previous methods get increasingly noisy with lags of increasing magnitude. Note that even though the correlation lasts on the order of a second or more (see vertical extension of central high correlation bands), estimates of the correlation at lags an order of magnitude smaller than that are quite noisy. This is due to the method of estimation as well as to any aperidocities that can exist in the signal.

---

1 If enough repetitions are available so as to allow the use of small windows, the length of the sliding window can be reduced even to a single sample with this method, for any lag range desired.

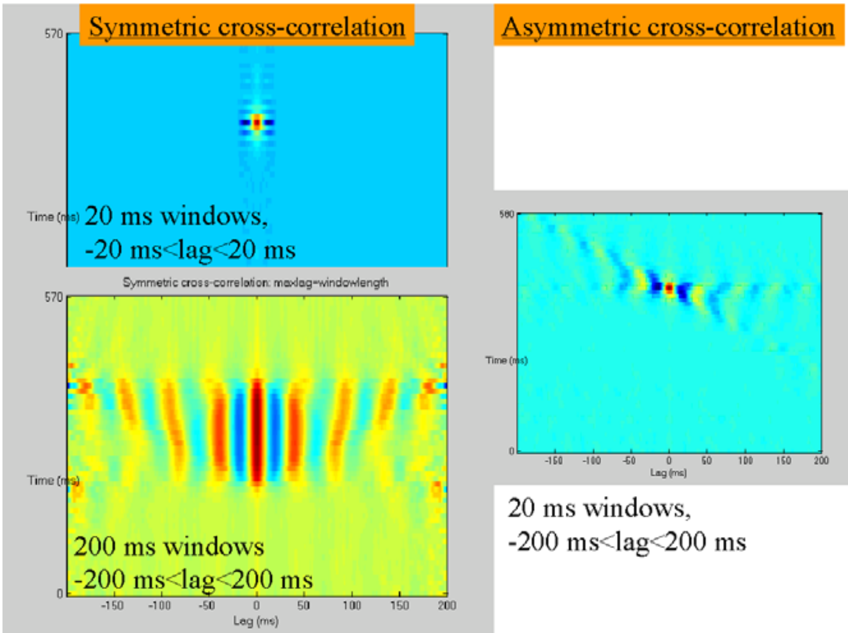

Figure 2. Symmetric cross-correlation forces a tradeoff between small windows, allowing for increased sensitivity to nonstationarities, and large lag ranges, allowing a full appreciation of any periodicity or delay in the correlation. Asymmetric cross-correlation allows for large lag ranges concurrent with small windows. The same signal is used for all 3 autocorrelograms above.